Argonne NATIONAL LABORATORY

Marching to Exascale: Extreme-Scale Cosmological Simulations with HACC* on Summit

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

CAAR

ECP EXASCALE COMPUTING PROJECT

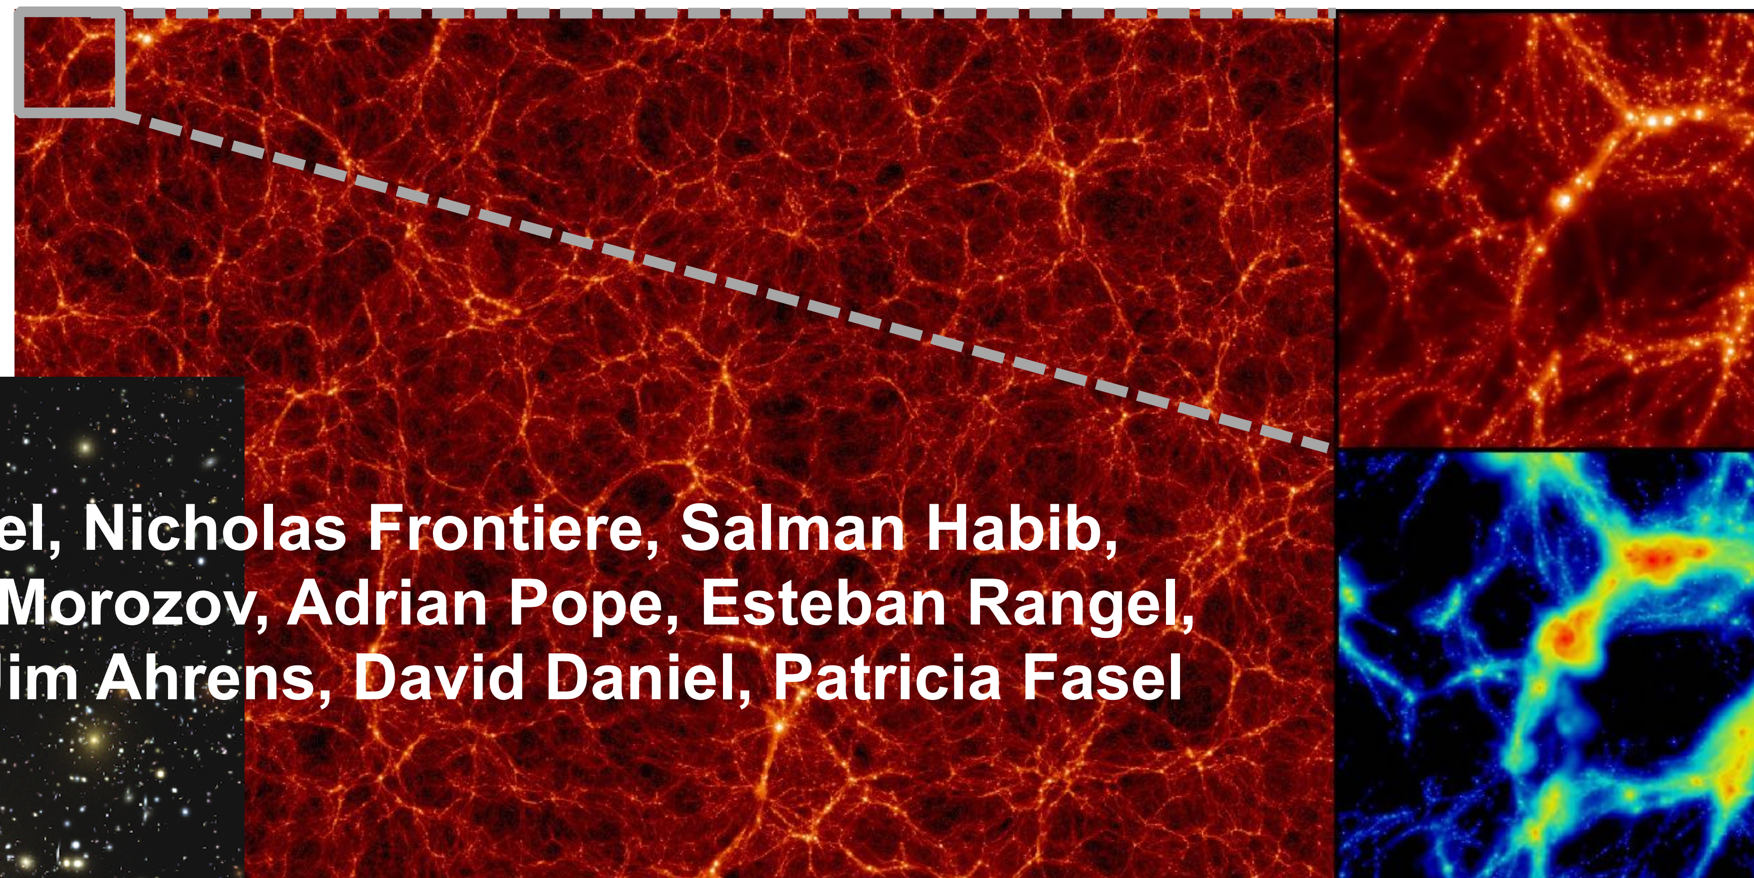

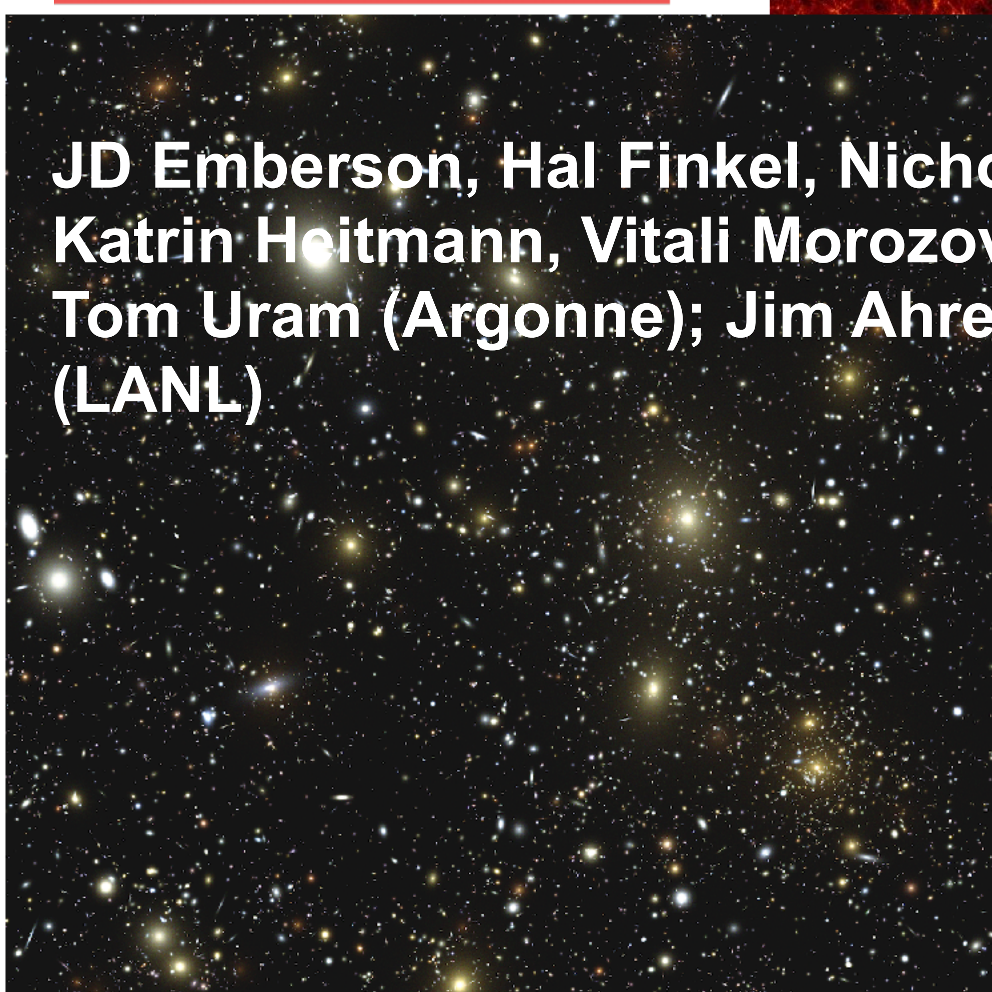

JD Emberson, Hal Finkel, Nicholas Frontiere, Salman Habib, Katrin Heitmann, Vitali Morozov, Adrian Pope, Esteban Rangel, Tom Uram (Argonne); Jim Ahrens, David Daniel, Patricia Fasel (LANL)

Salman Habib
Argonne National Laboratory

2019 OLCF User Meeting
May 21, 2019

*Hardware/Hybrid Accelerated Cosmology Code

A Little Personal History —

- **Supercomputer Evolution:** Supercomputing has come a very long way since summer 1993 when the CM-5 was #1 —
  - 1024 nodes, 131 GFlops (peak)! (131 kW)
  - Total RAM, 32 GB!
  - Could barely do 512^3 simulations
  - But it was easy and fun to program!

- **Summit is a world apart:** No more 1-2 people codes, now have end-to-end simulations, etc.
  - Focus on code architecture and algorithms
  - Complex science problems
  - Effective teaming essential
  - System-scale modeling/simulation
  - Still fun (in a way) —

HACC Science Drivers: Cosmological Surveys

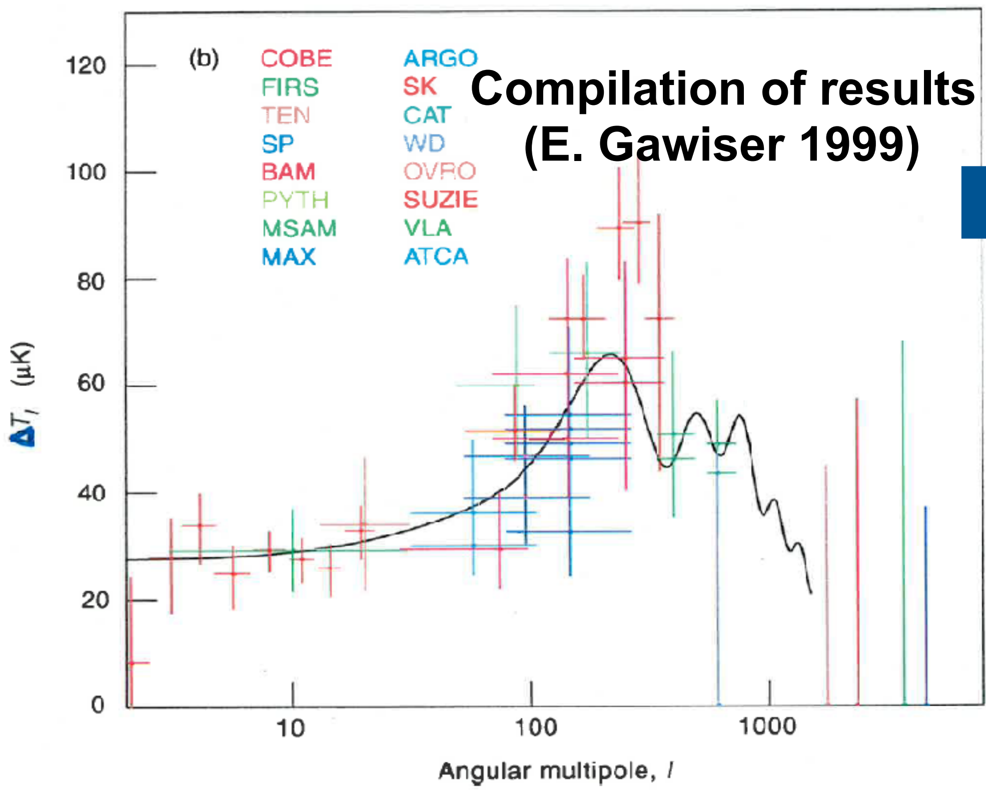

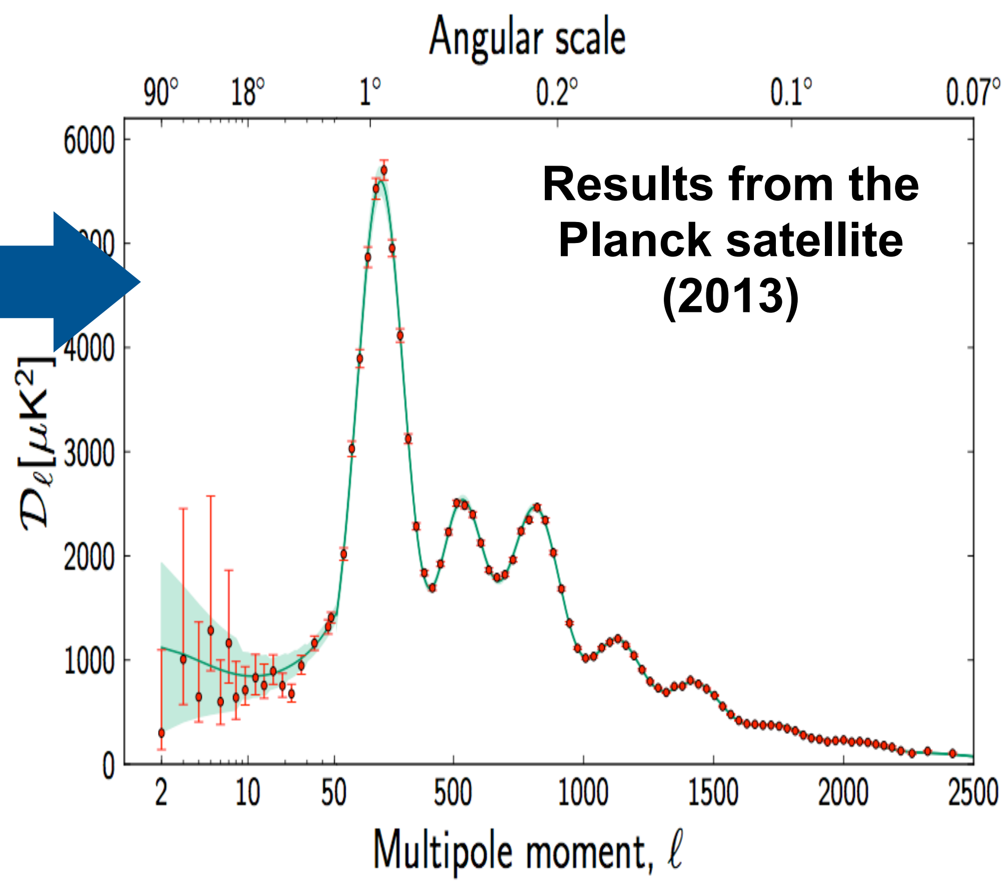

- Massive increase in sensitivity of cosmic microwave background (CMB) observations
- Cross-correlation with galaxy surveys
- New era of CMB modeling/simulations

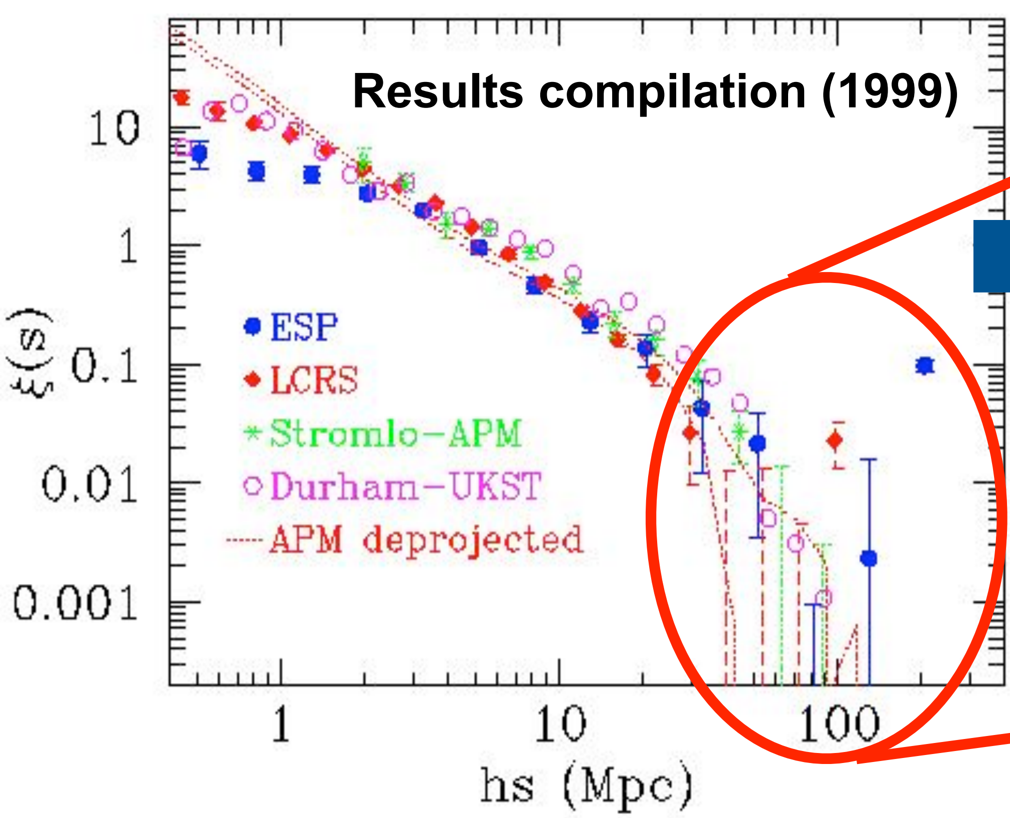

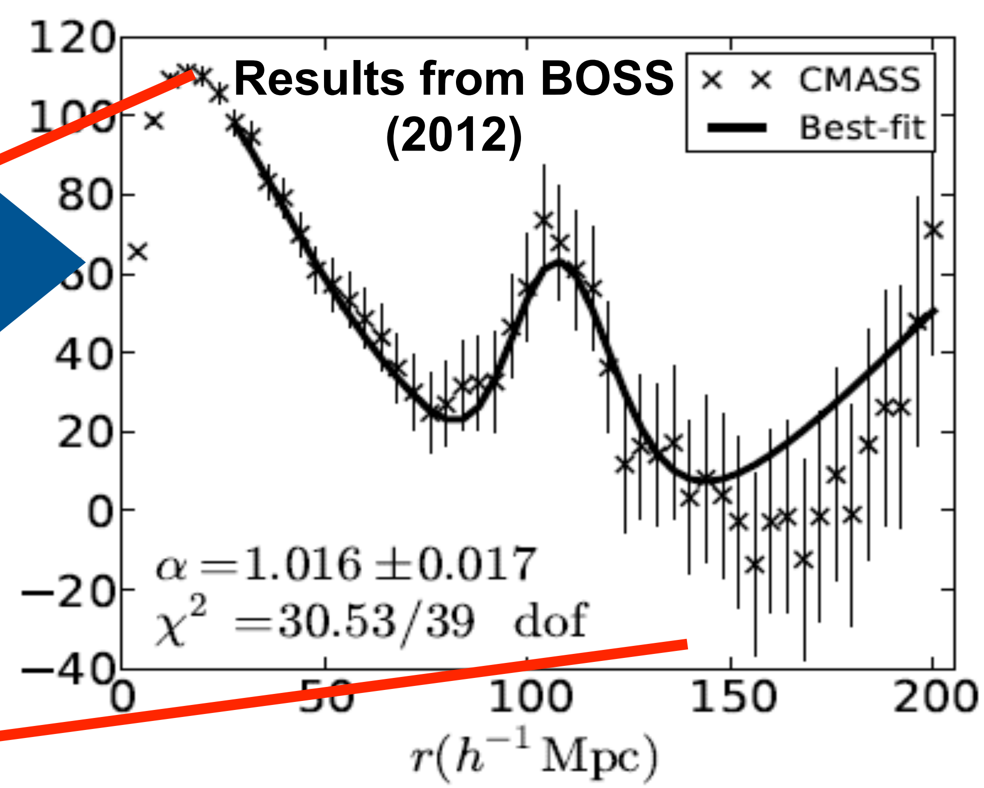

- Massive increase in volume of galaxy surveys
- Next-generation galaxy clustering simulations
- Multi-physics codes needed to meet accuracy requirements

Cosmology: Simulation Frontiers

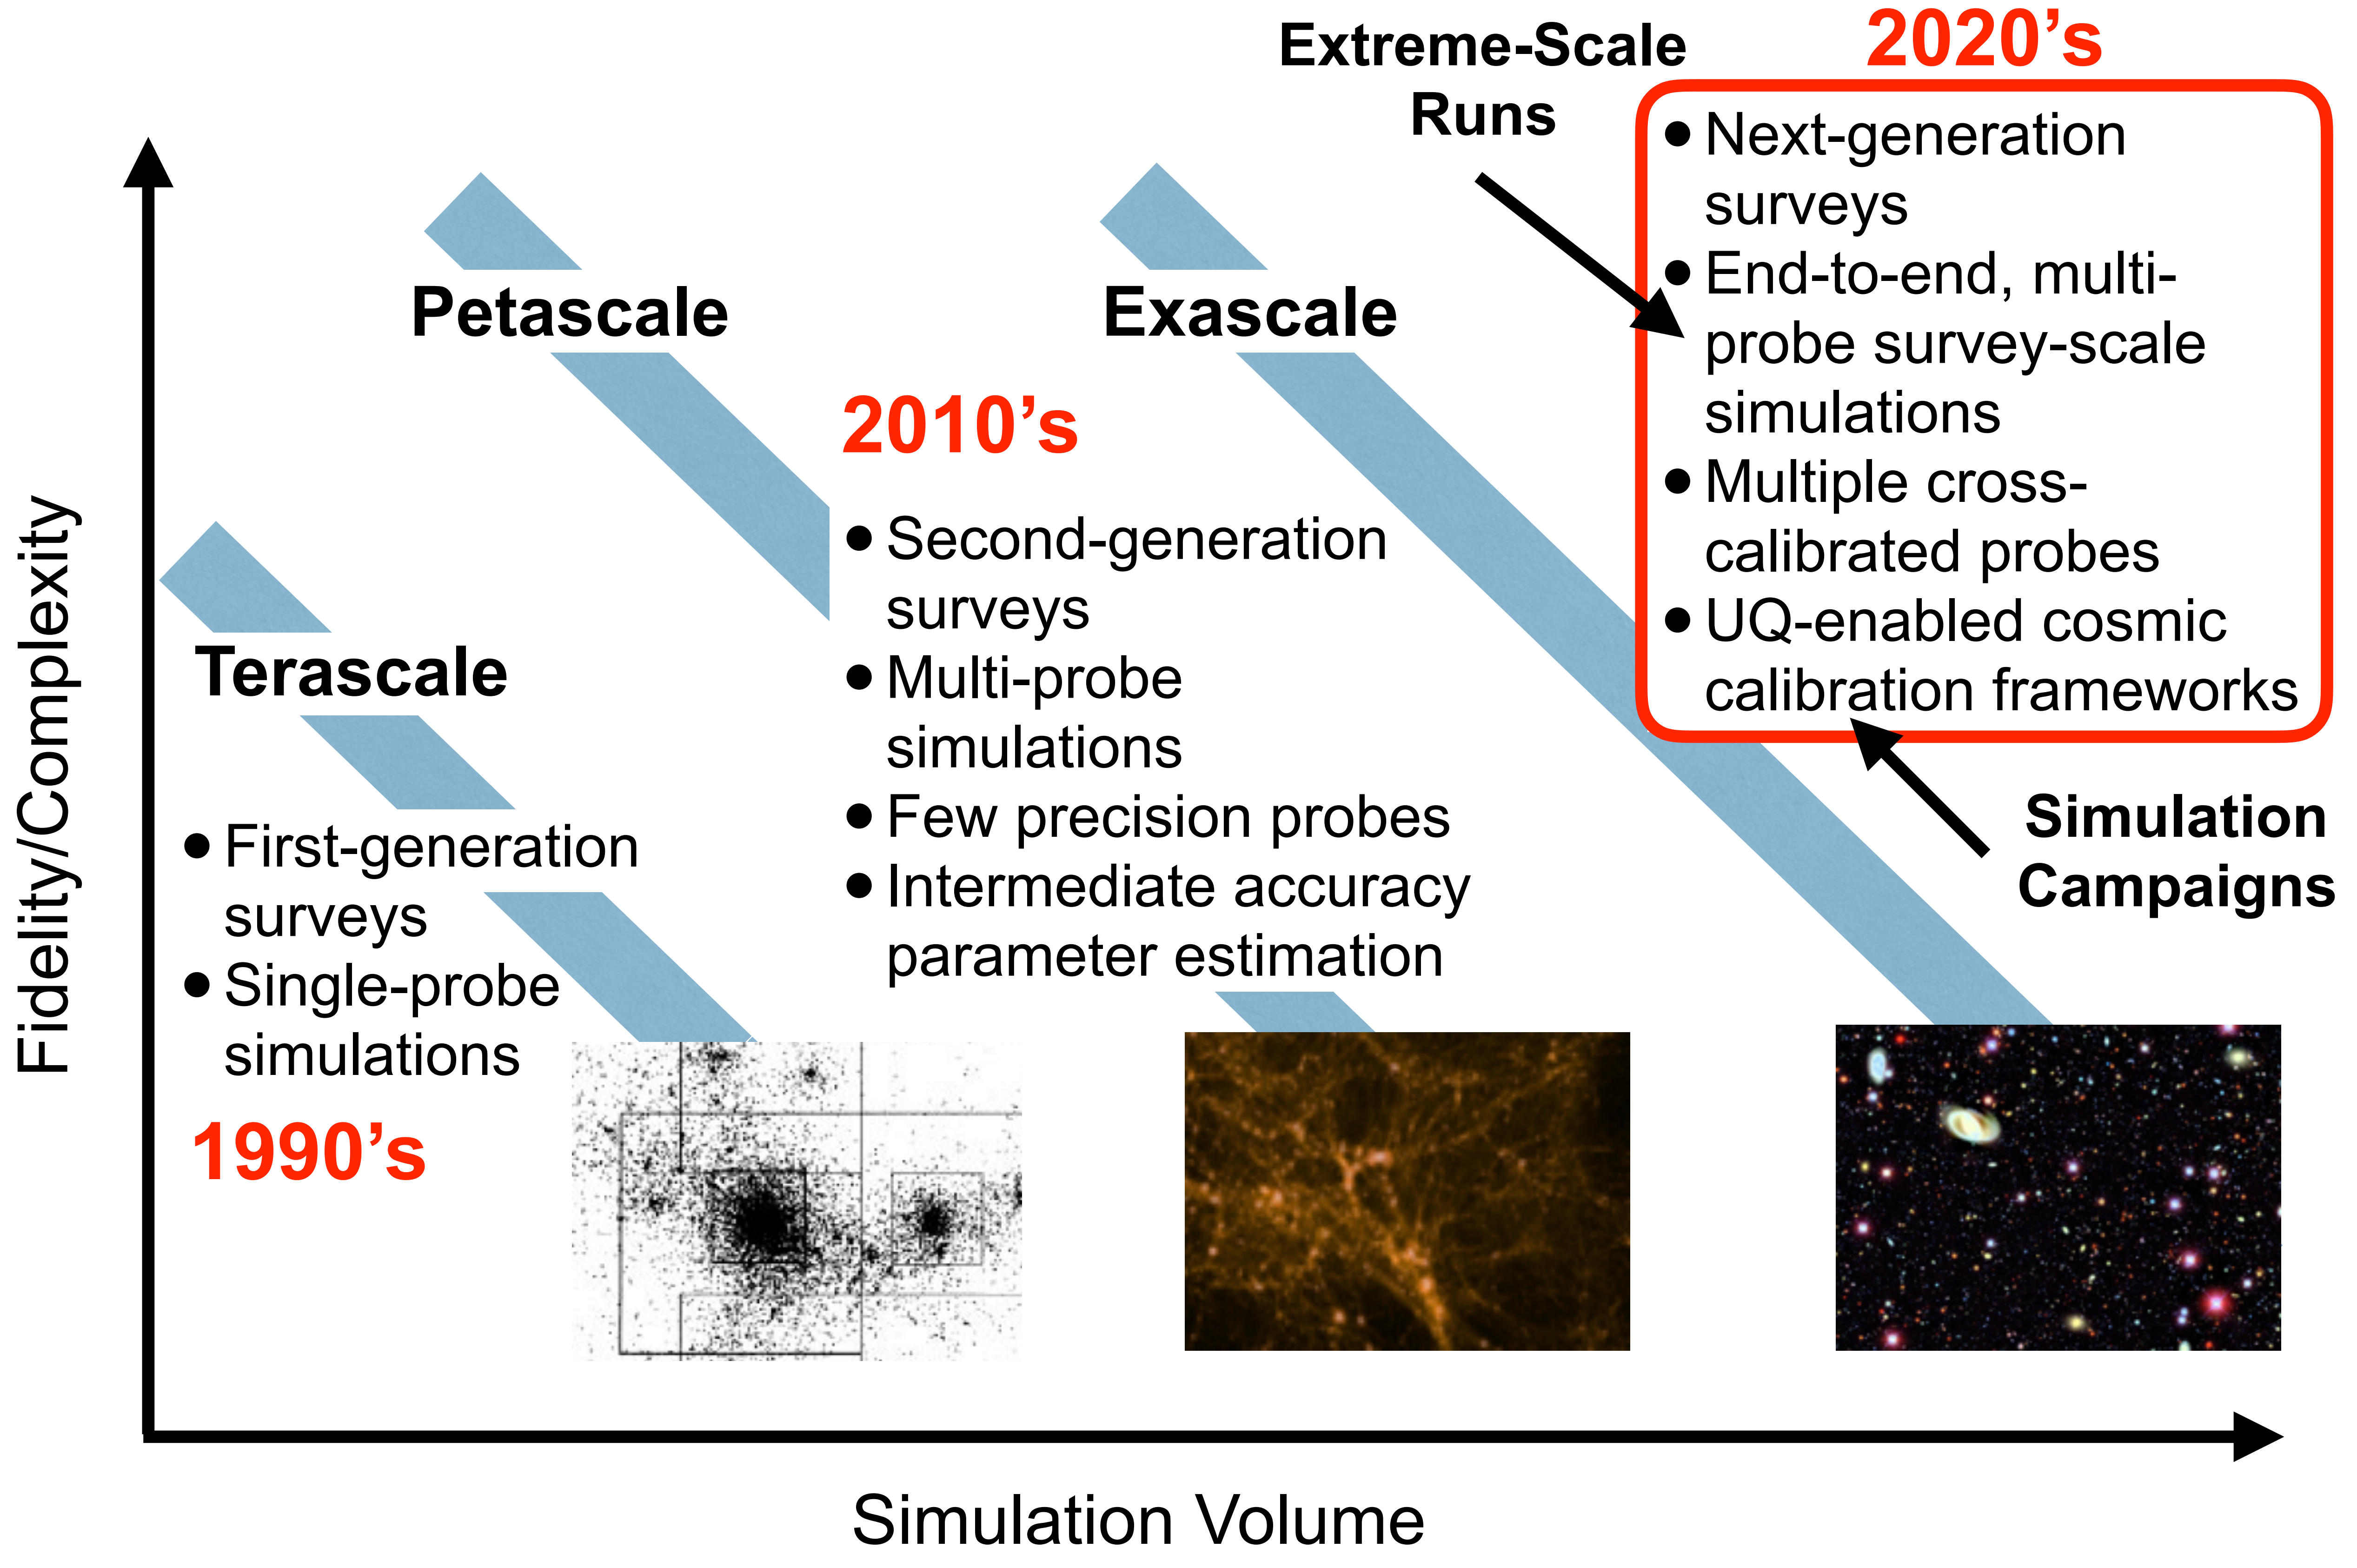

End-To-End Workflow Complexity

Gaussian Random Field Initial Conditions → High-Resolution N-Body (+ Hydro) simulations → Multiple Outputs Halo/Sub-Halo Identification → Halo Merger Trees

HACC

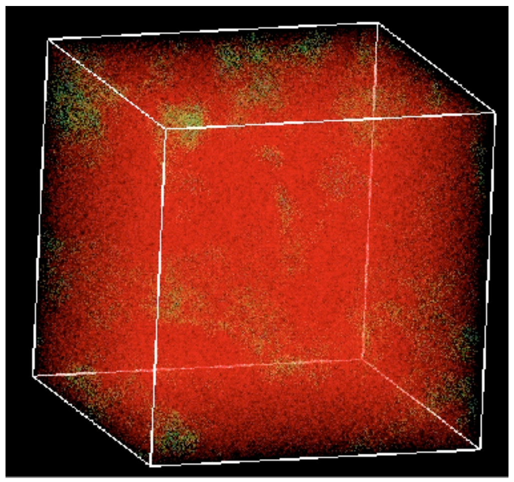

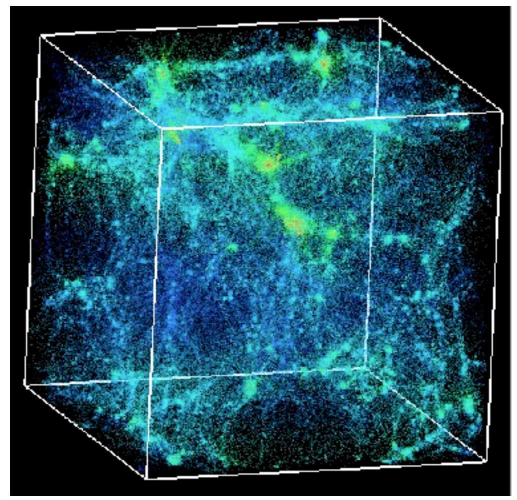

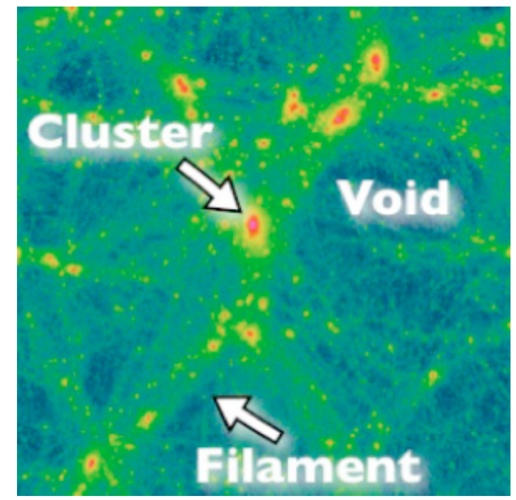

Halo Merger Trees → Semi-Analytic + Empirical Galaxy Modeling → Galaxy Catalog → Realistic Image Catalog → Atmosphere and Instrument Modeling

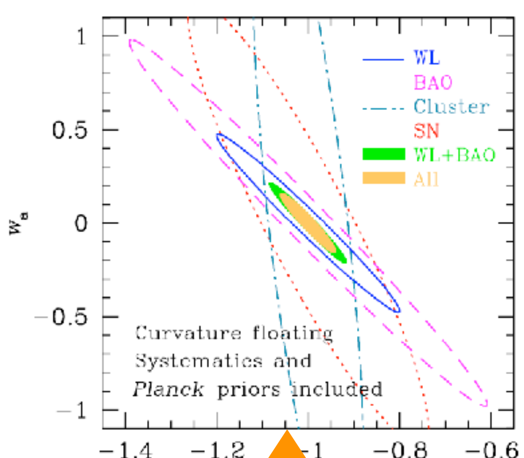

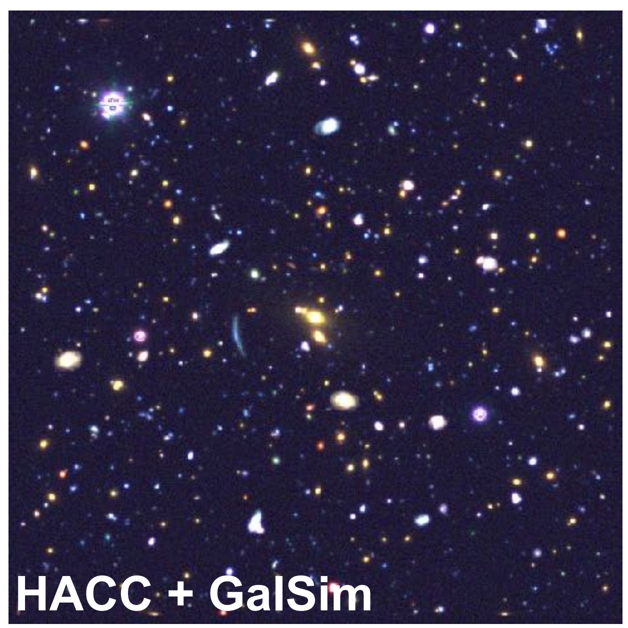

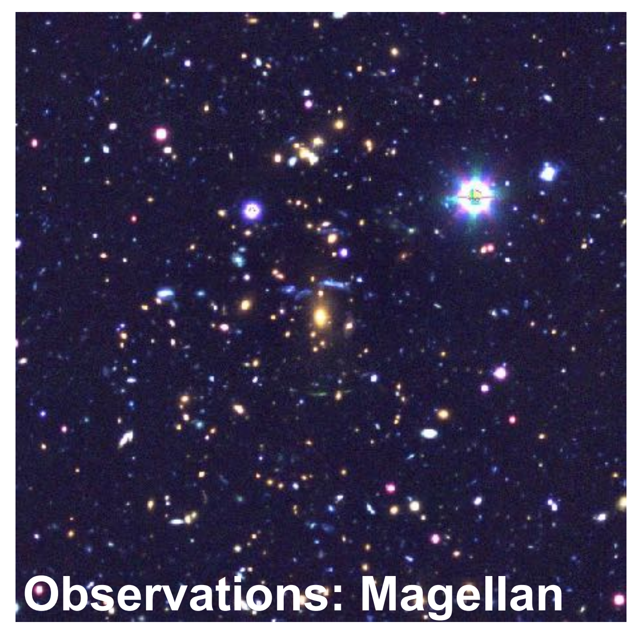

Atmosphere and Instrument Modeling → Data Management Pipeline → Data Analysis Pipeline → Scientific Inference Framework

Cosmological Inverse Problem

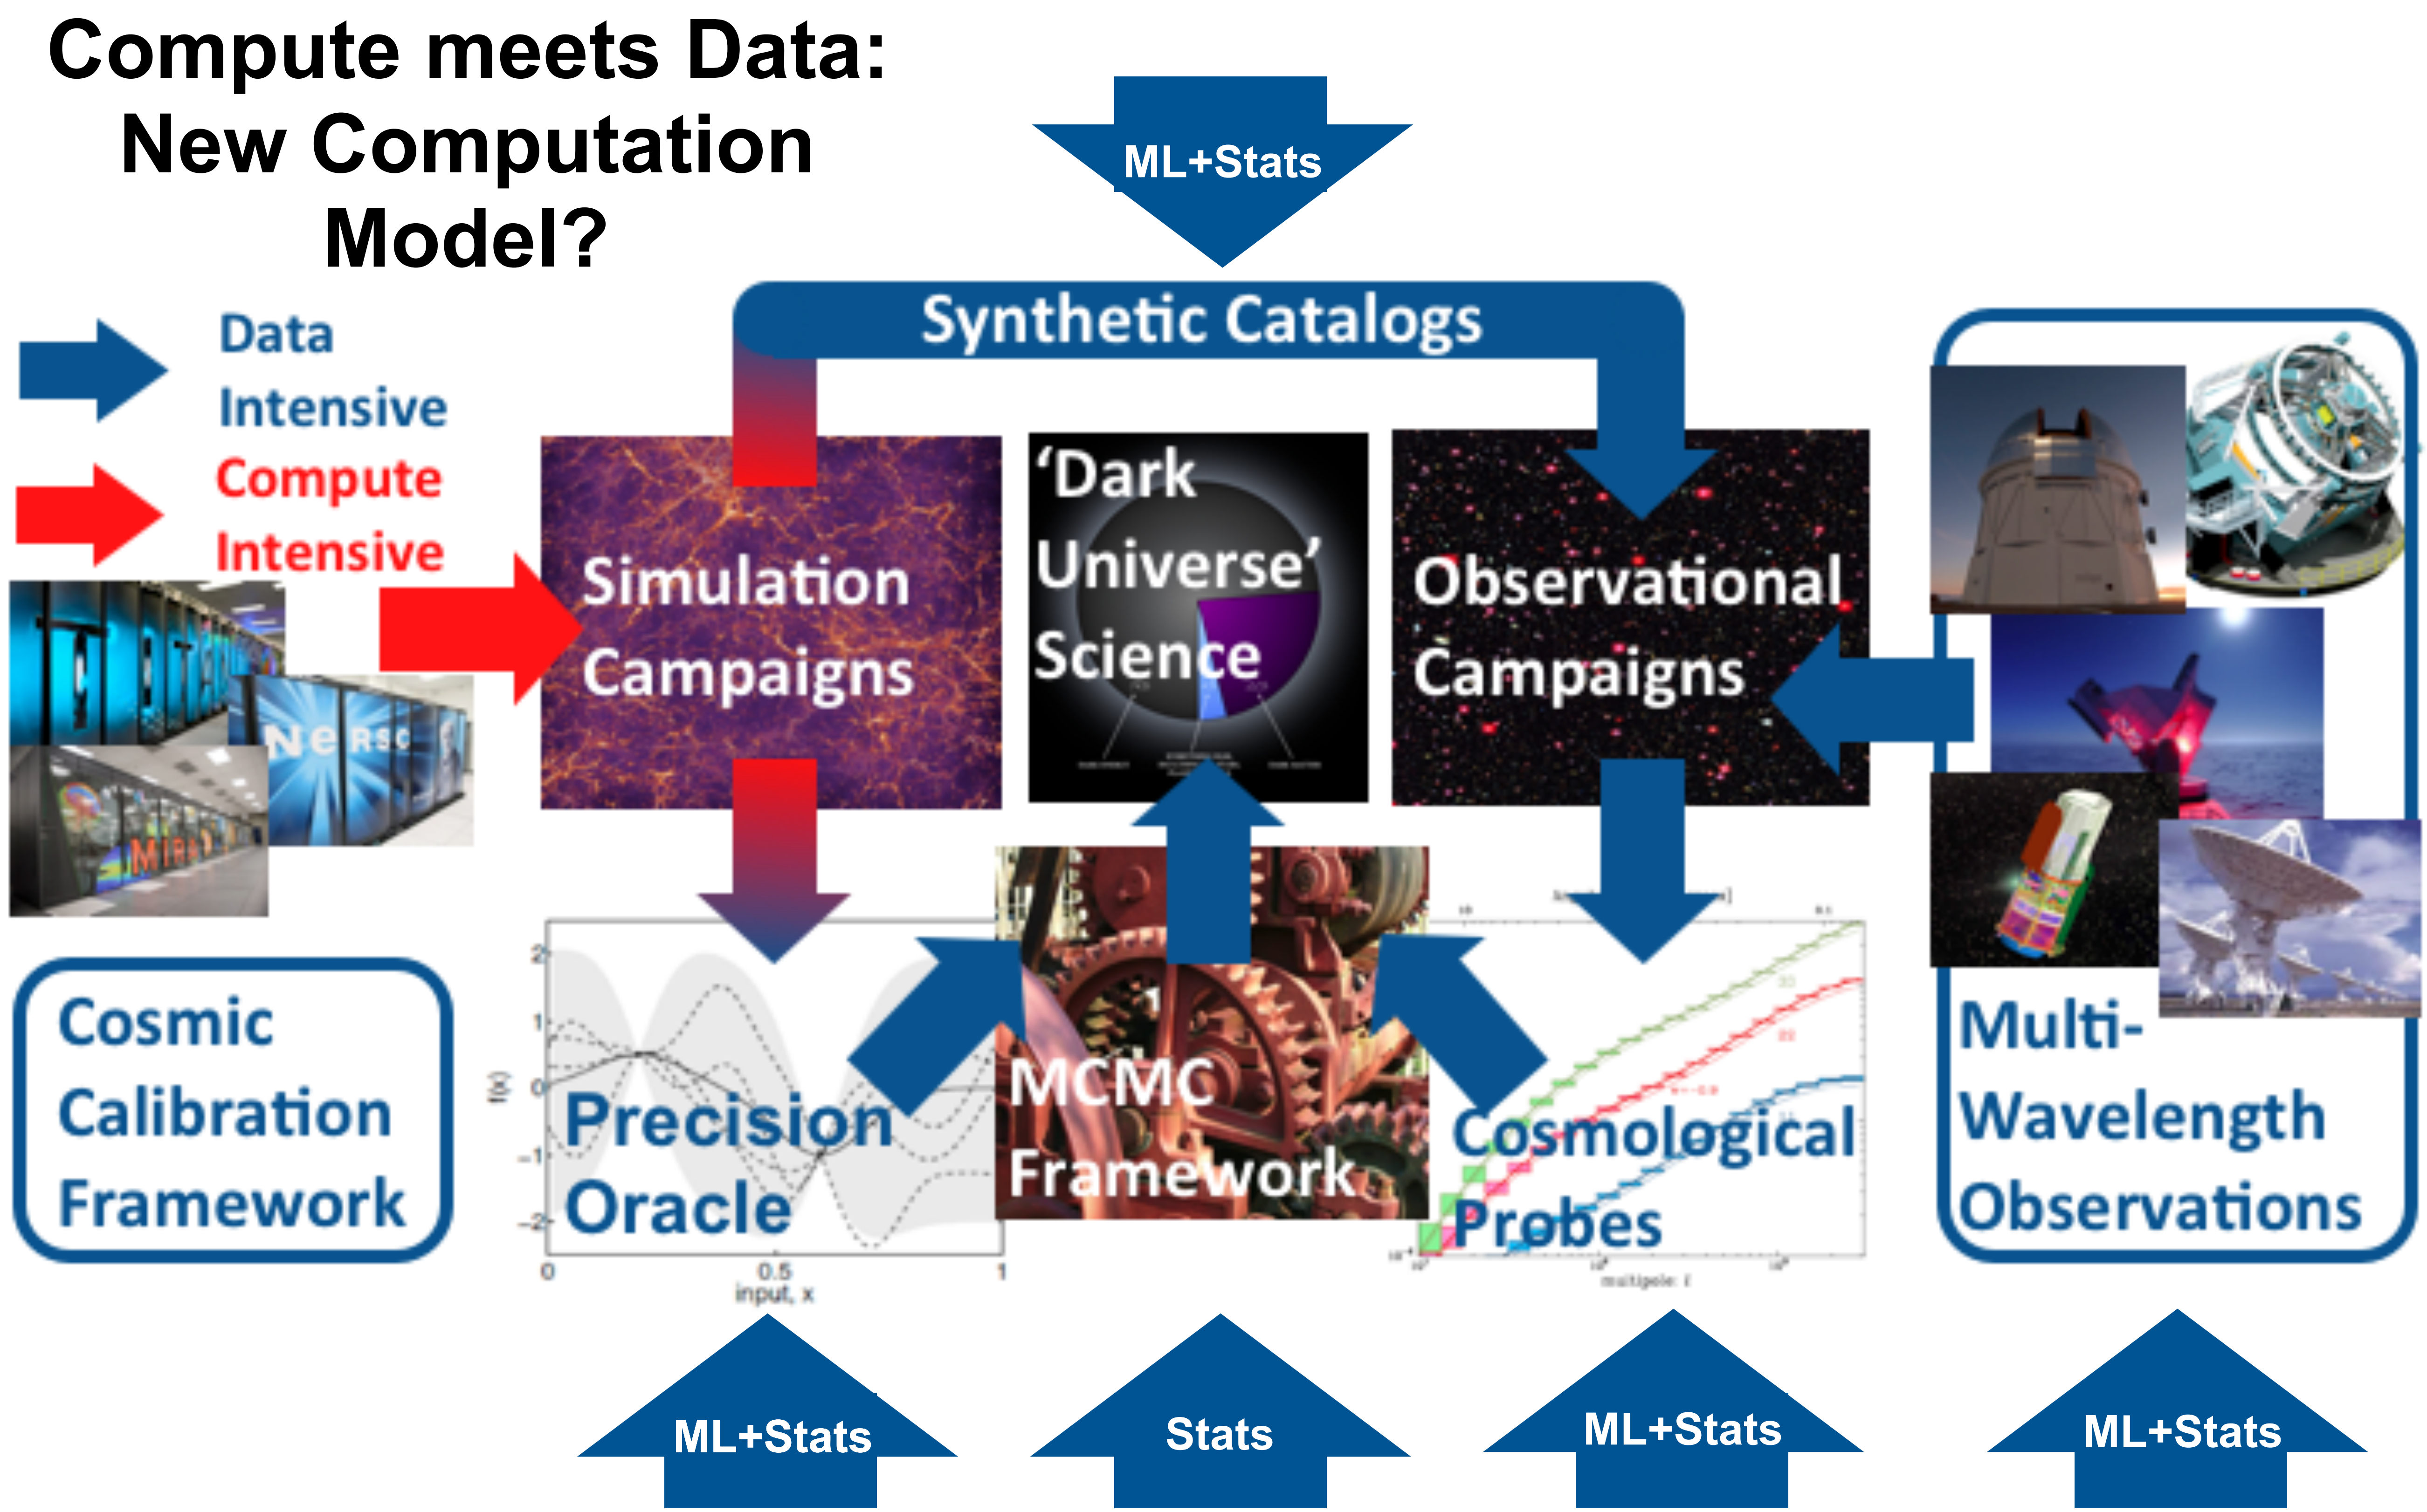

Science Drivers for Future HACC Development

- **Science Requirements:** Next-generation survey science synthetic catalogs; detailed modeling for cosmological probes and cross-correlations; large-scale systematics studies
  - Weak lensing (WL) shear
  - Galaxy clustering (LSS)
  - Multiple LSS X WL cross-correlations
  - Cluster cosmology (mass calibration)
  - Secondary CMB anisotropies and backgrounds
  - Multiple CMB X LSS cross-correlations

Many billions of objects, ~500 parameters each

Synthetic extra-galactic image from a HACC simulation (for LSST DESC)

LSST

>1000 member collaborations breathing down our necks!

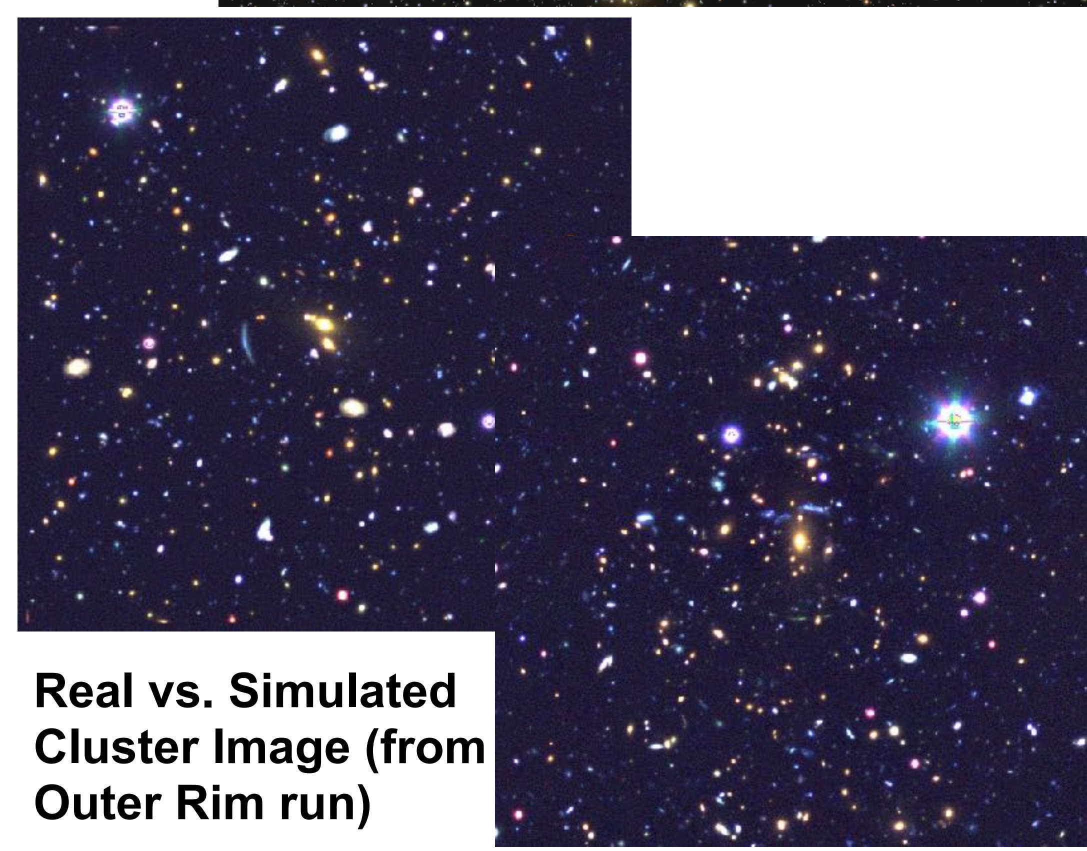

Real vs. Simulated Cluster Image (from Outer Rim run)

HACC (Hardware/Hybrid Acelerated Cosmology Code)

- **HACC Physics and Problem Scale:** extreme-scale, particle-based (Lagrangian) framework for computational cosmology; solves the Vlasov-Poisson equation using particles, includes gas dynamics (new algorithm: CRK-SPH), with astrophysical feedback

Frontiere et al., JCP 332, 160 (2017)

- **HACC Design Priorities:**
  - **Very high levels of performance** across multiple architectures — focus on ***absolute,*** not just ***relative*** performance, ~50%+ of peak
  - Be **fully scalable** (strong/weak scaling on the largest platforms)
  - **Performance portability:** uses multiple algorithms and implementations across architectures, *yet 95% of the code base remains unchanged*
  - Run on **pre-deployment systems** as a benchmark
- **Programming Model:** C++/MPI + X (X = OpenMP/ CUDA/OpenCL/assembly)

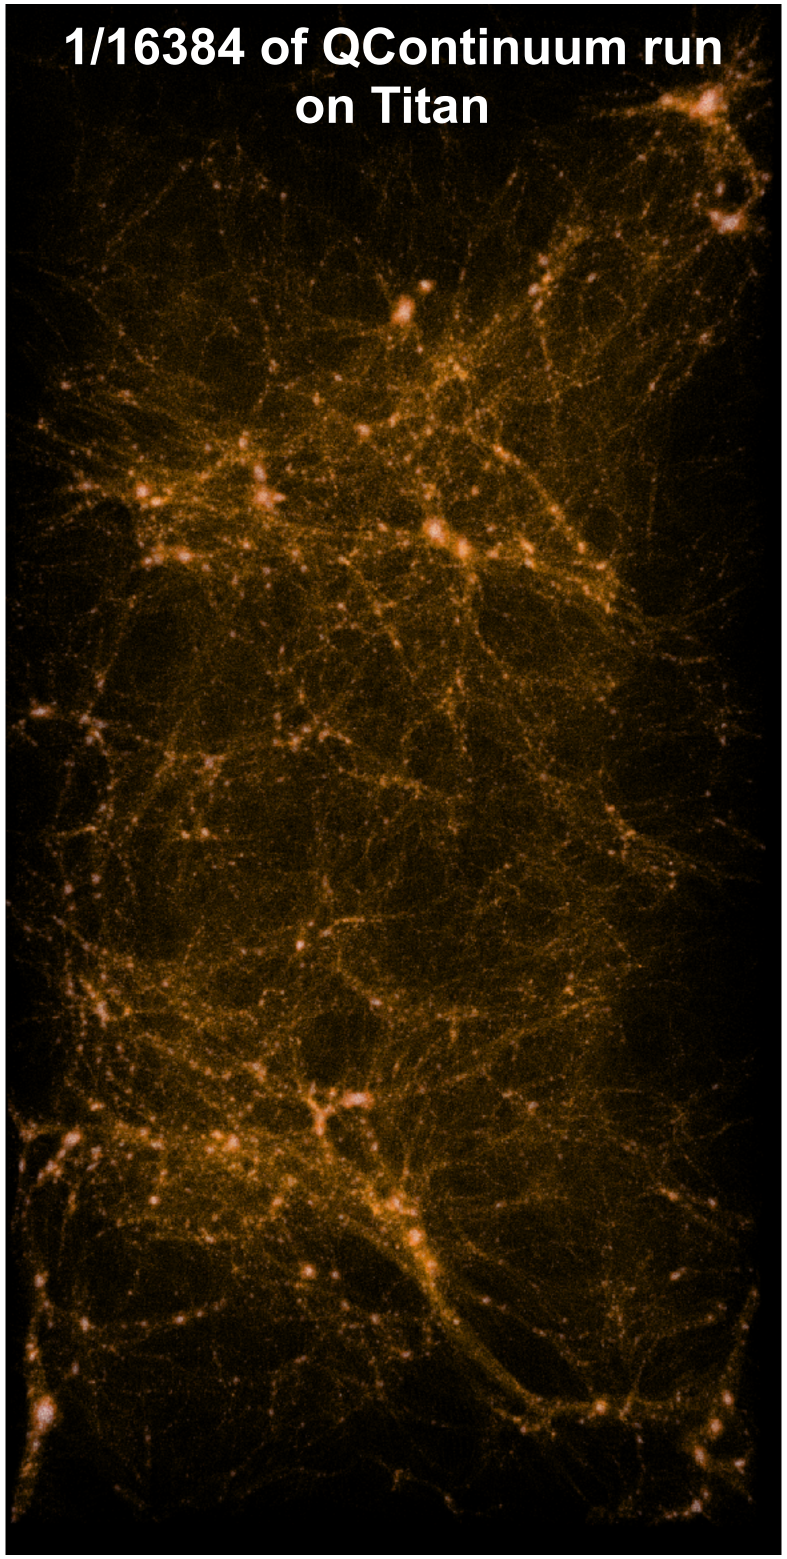

Habib et al., New Astron. 42, 49 (2016); Comm. ACM 60, 97 (2017) [Research Highlight]

HACC: Algorithmic Features

- **Philosophy:** Smooth physics (large scale) — grids (architecture invariant); small-scale "rough" physics — particles (architecture-tuned)
- **Gravity — Hybrid Grid/Particle:** 6-th order spectral Poisson solver; 4-th order super-Lanczos spectral derivatives; short-range forces via spectral filters (high-accuracy polynomial fits), custom parallel 3D FFT
- **Gasdynamics — CRK-SPH:** Higher-order SPH scheme solves known SPH issues in dealing with mixing, tracking instabilities, etc.
- **Flexible Chaining Mesh and Local Trees:** Data structures optimize local force solvers (tree/fast multipole/P3M); neighbor list computation
- **Adaptive (Symplectic) Time-Stepping:** 2nd-order split-operator method; sub-cycling based on the RCB tree depth; implicit solver for subgrid models

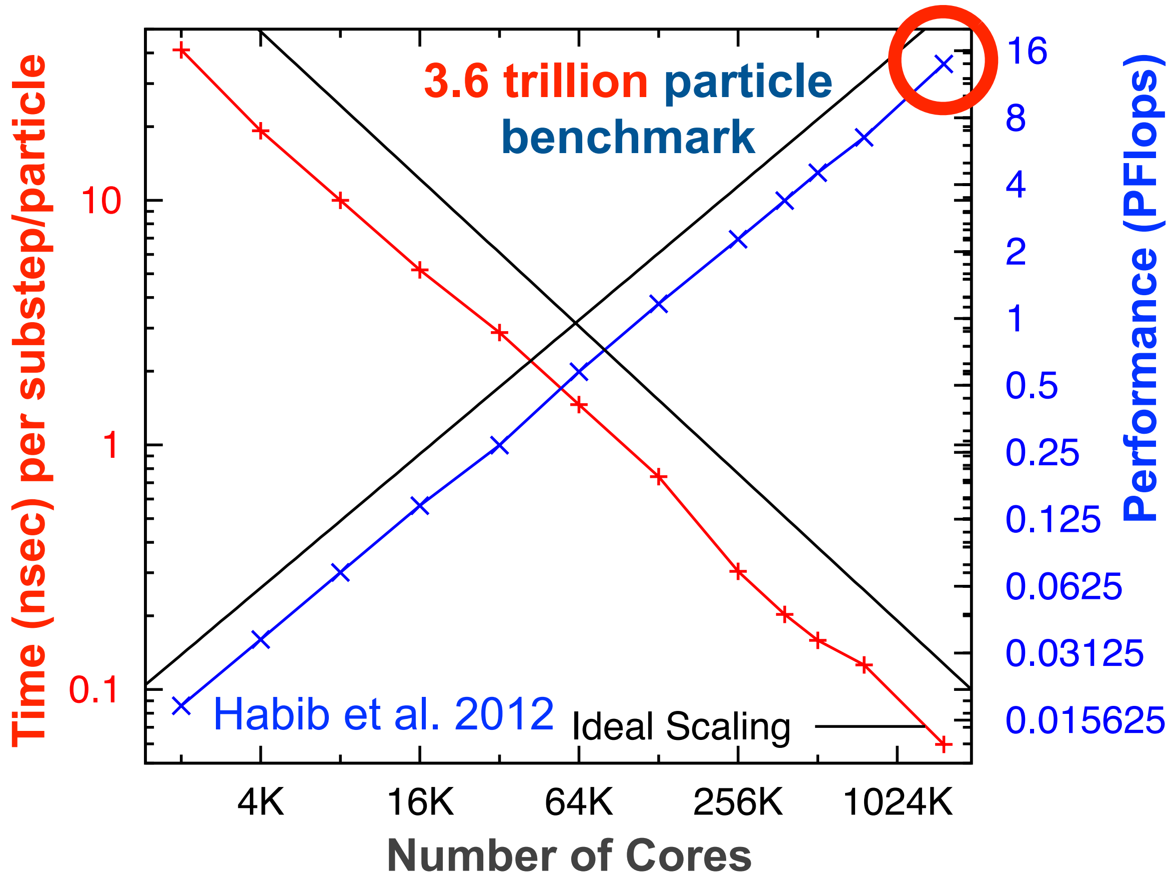

HACC on Sequoia (2012): 13.94 PFlops, 69.2% peak, 90% parallel efficiency on 1,572,864 cores/MPI ranks, 6.3M-way concurrency

HACC In Pictures (“Million to One” Dynamic Range)

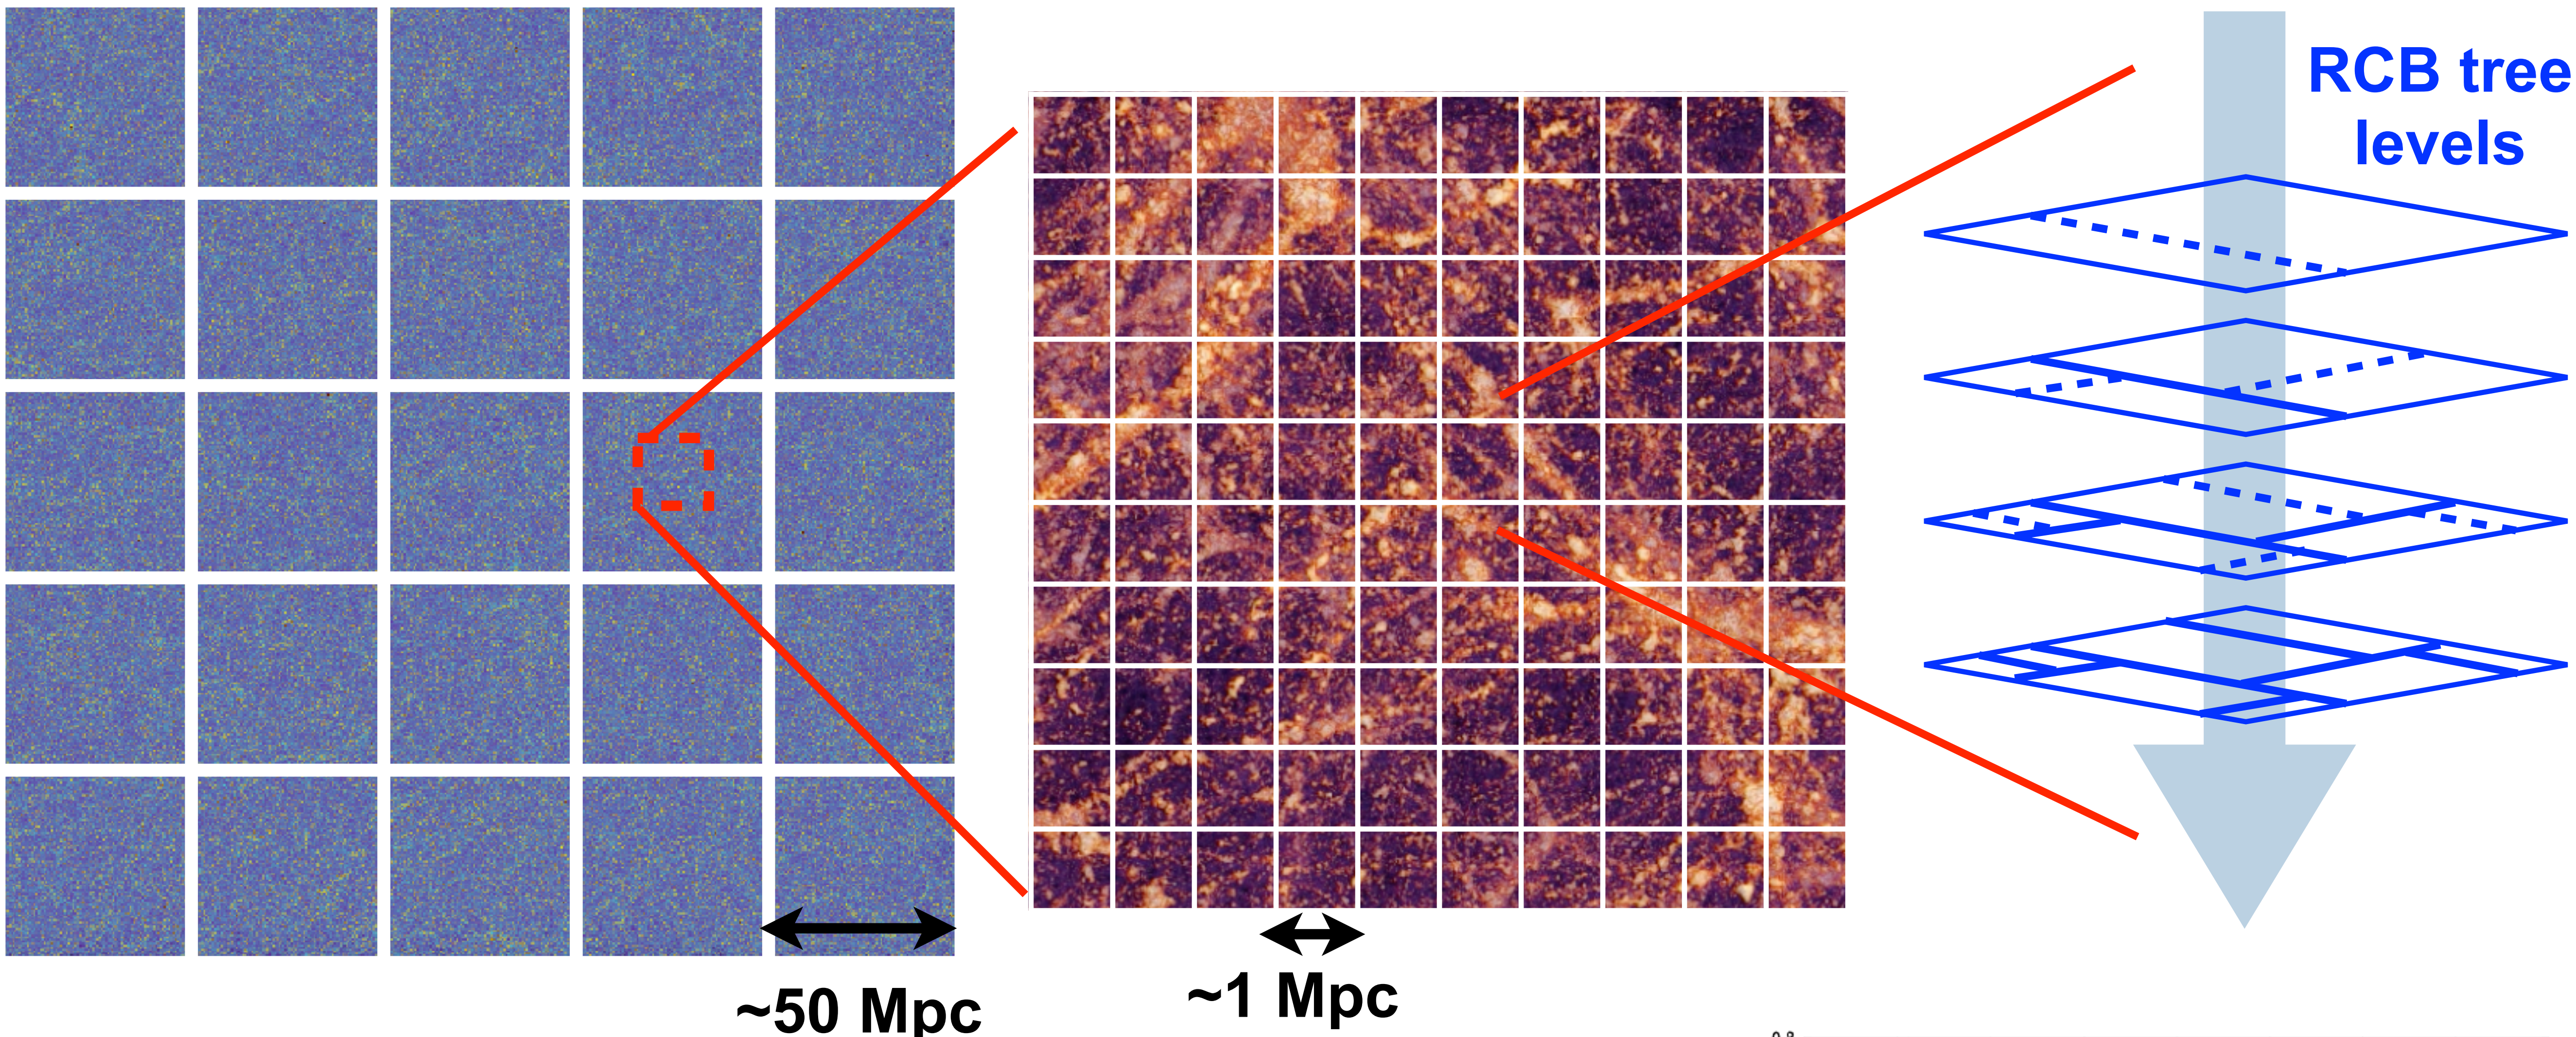

HACC Top Layer:

3-D domain decomposition with particle replication at boundaries (‘overloading’) for Spectral PM algorithm (long-range force)

Host-side

HACC ‘Nodal’ Layer:

Short-range solvers employing combination of flexible chaining mesh and RCB tree-based force evaluations

GPU (compute-intensive): two options, P3M vs. TreePM

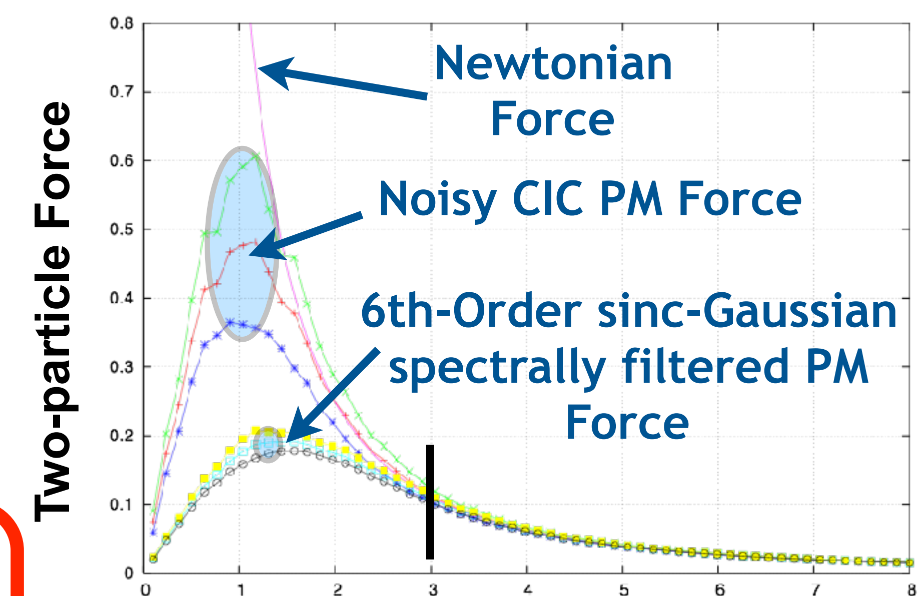

HACC on Summit Phase 1/2

- **Transition from Titan:** In cosmological simulations, the computational intensity increases as a function of time due to increased particle clustering, this was a problem on Titan (and even worse on manycore systems)
- **Initial Runs on Summit Phase 1/2:** HACC ran on Summit within hours of the machine being available — detailed timing tests show expected weak and strong scaling results

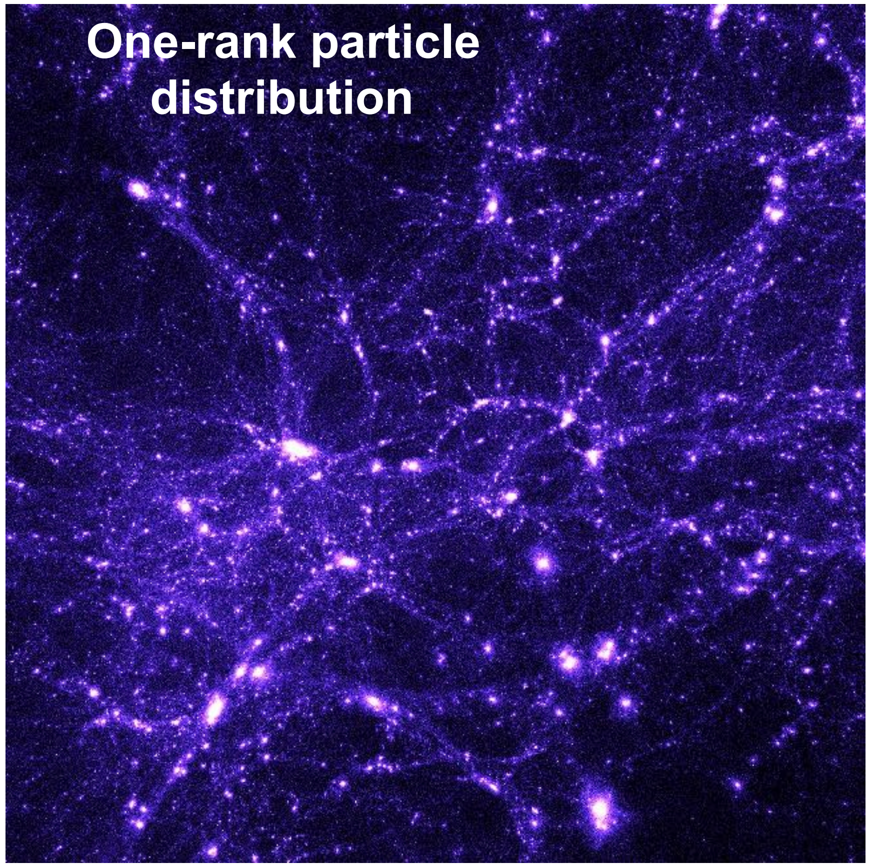

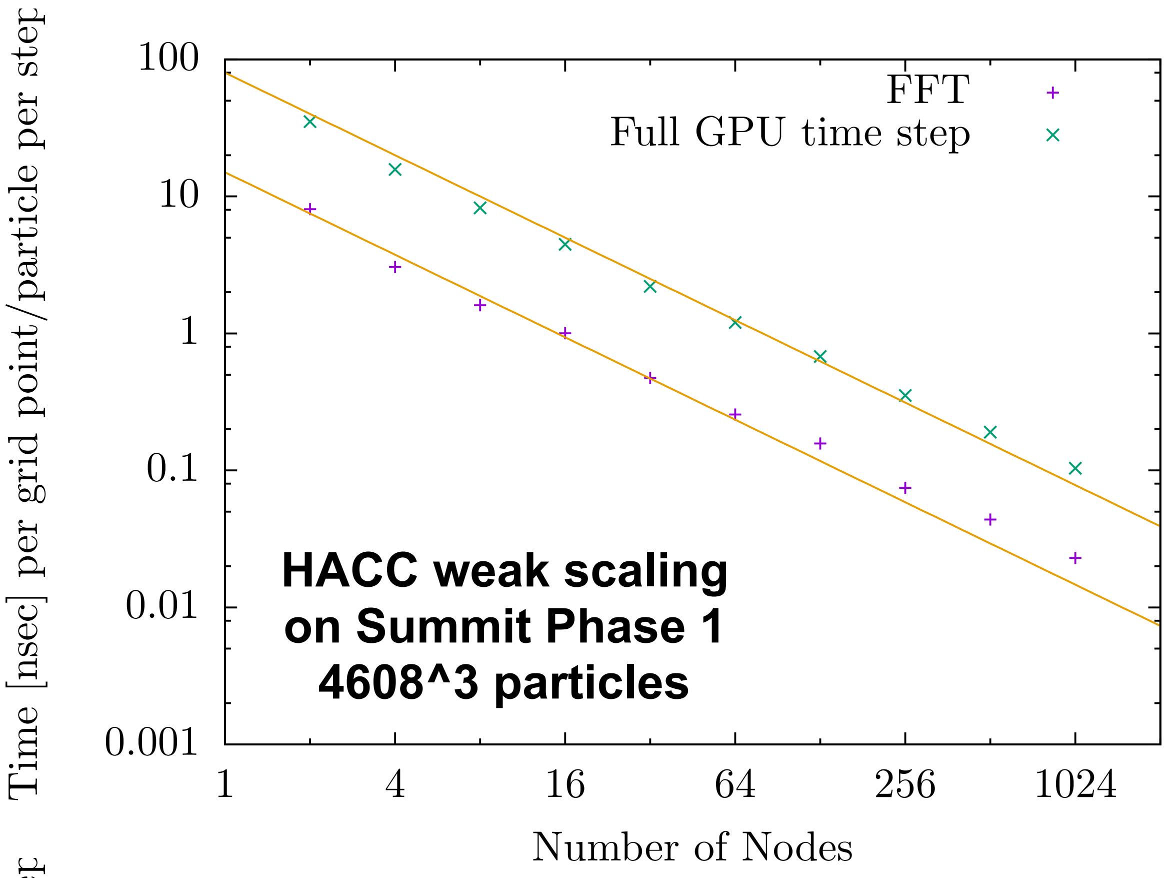

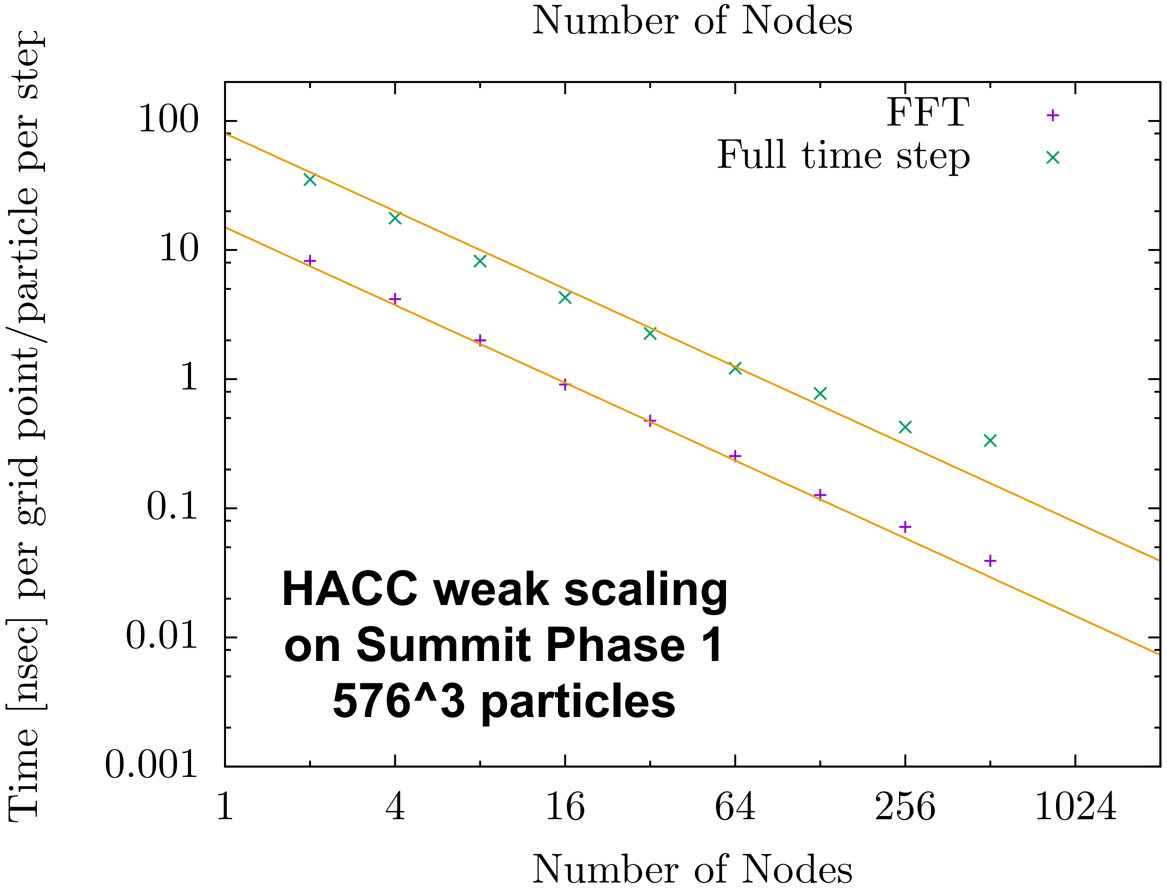

HACC on Summit I

- **Major Runs:** Carried out on Summit during acceptance period, 14 runs each set up for 24 hours, 9 finished cleanly; 3 of the world's largest cosmological simuations consumed roughly a week of machine time
- **Slow-Down Problem Solved:** Only ~30% slowdown as a function of time, with full timesteps taking only ~5 minutes!

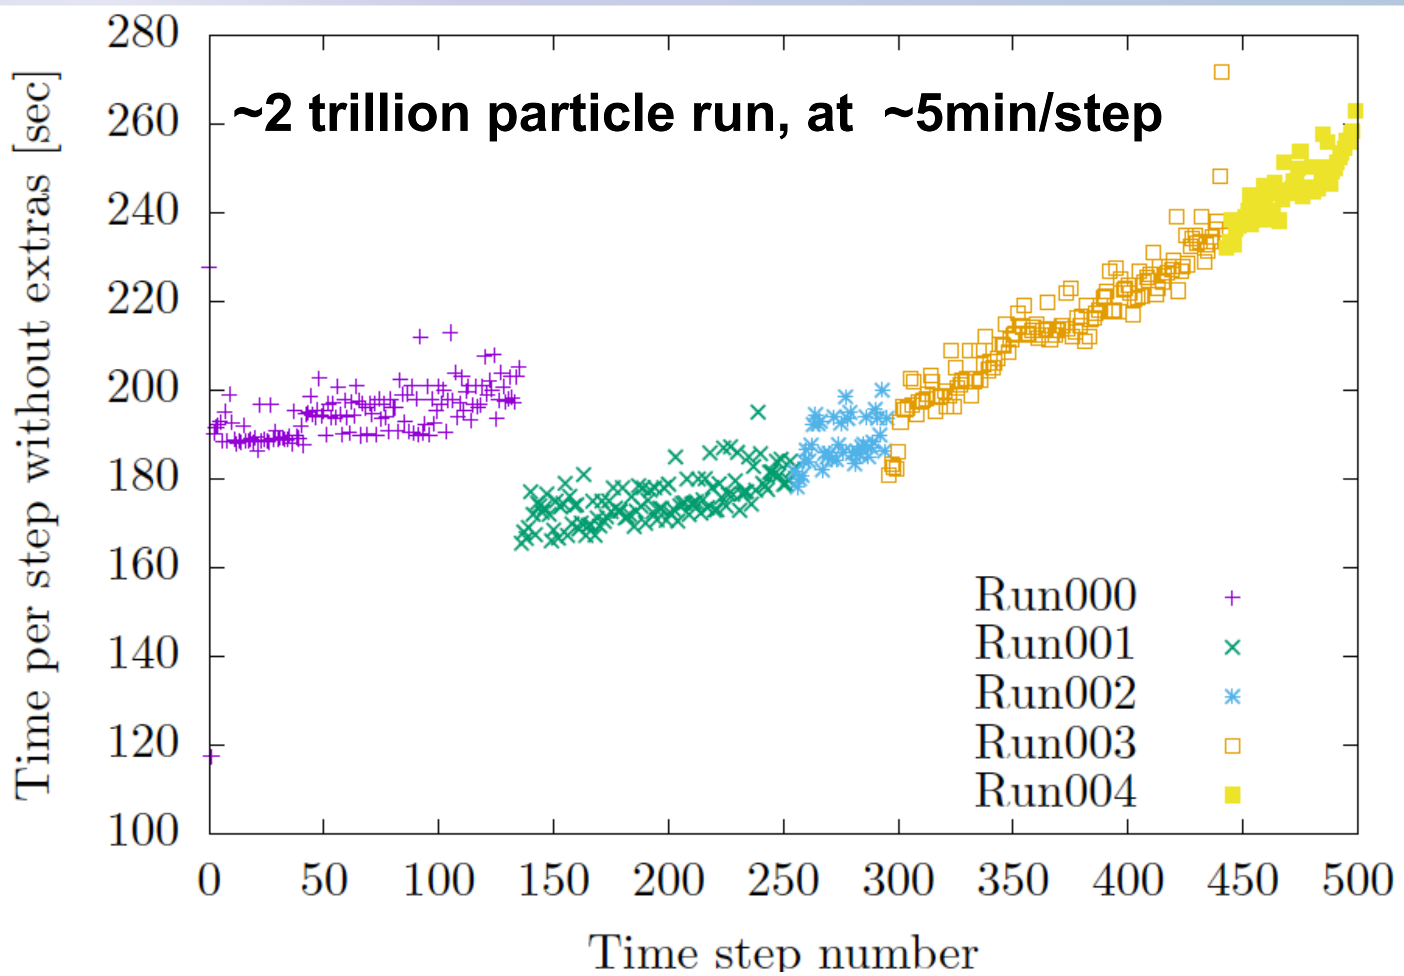

Simulation	Run	Job #	Exit	Error Message	Problem
Qo'noS	Run003a	229161	hang	Could not read jskill result from pmix server	halo finder particle exchange
Qo'noS	Run000	229198	clean	n/a	
Qo'noS	Run001	229290	clean	n/a	
Qo'noS	Run002	229467	crash	error 12	FFT
Qo'noS	Run003	229494	clean	n/a	
Qo'noS	Run004	229618	clean	n/a	
Vulcan	Run000	230003	hang	Could not read jskill result from pmix server	Halo finder i/o(?) after *Total halo particles*
Vulcan	Run001	230122	clean	n/a	
Vulcan	Run002	230178	clean	n/a	
Vulcan	Run003	230304	crash	Segmentation fault	after short range solver, no *Allocated heap:*
Vulcan	Run004	230337	clean	n/a	
Ferenginar	Run000	230380	clean	n/a	
Ferenginar	Run001	231468	crash	Segmentation fault	in halo center finder
Ferenginar	Run002	231911	crash	error 12	in halo center finder
Ferenginar	Run003	232210	clean	n/a	

HACC on Summit II

- **Simulation Specs:** Largest simulations at a mass resolution of 10^9 M_sun
  - Spatial dynamic range: million to one (3Gpc box, ~3kpc force resolution)
  - ~2 trillion particles
  - 4096 Summit nodes
  - 3 cosmologies to distinguish subtle effects
  - **Qo'noS:** best-fit Planck LCDM cosmology
  - **Ferenginar:** early dark energy model + Planck
  - **Vulcan:** small massive neutrino component (0.1eV mass) + Planck
  - Large analysis suite for multiple cosmological probes: weak/strong gravitational lensing, galaxy clustering, galaxy clusters, —
  - Total data generated ~9PB (should keep analysis teams busy for years)

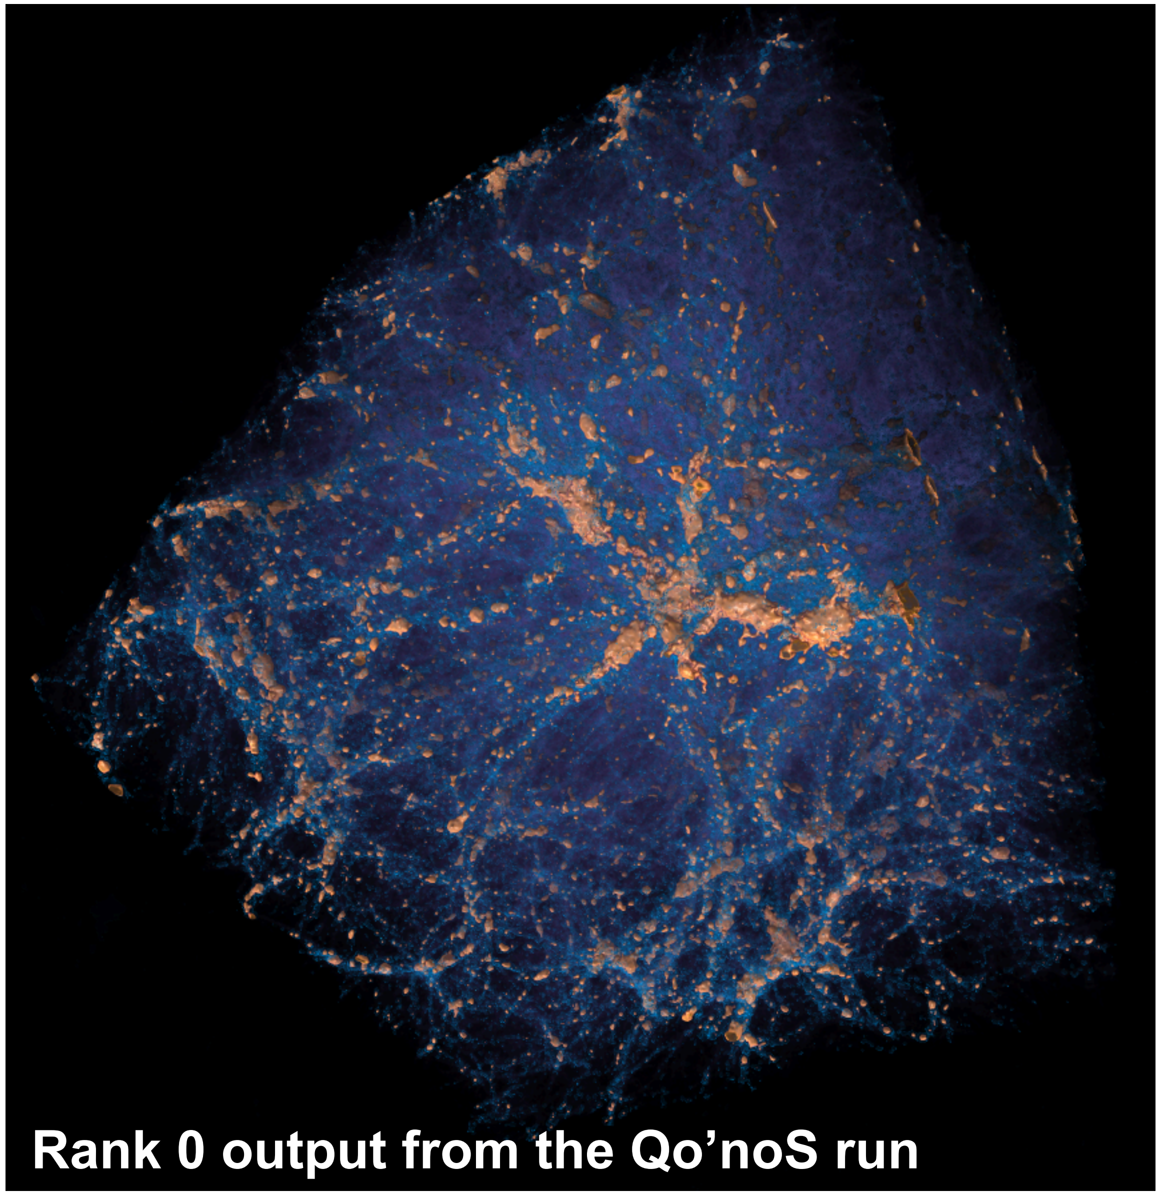

Rank 0 output from the Qo'noS run

HACC on Summit III

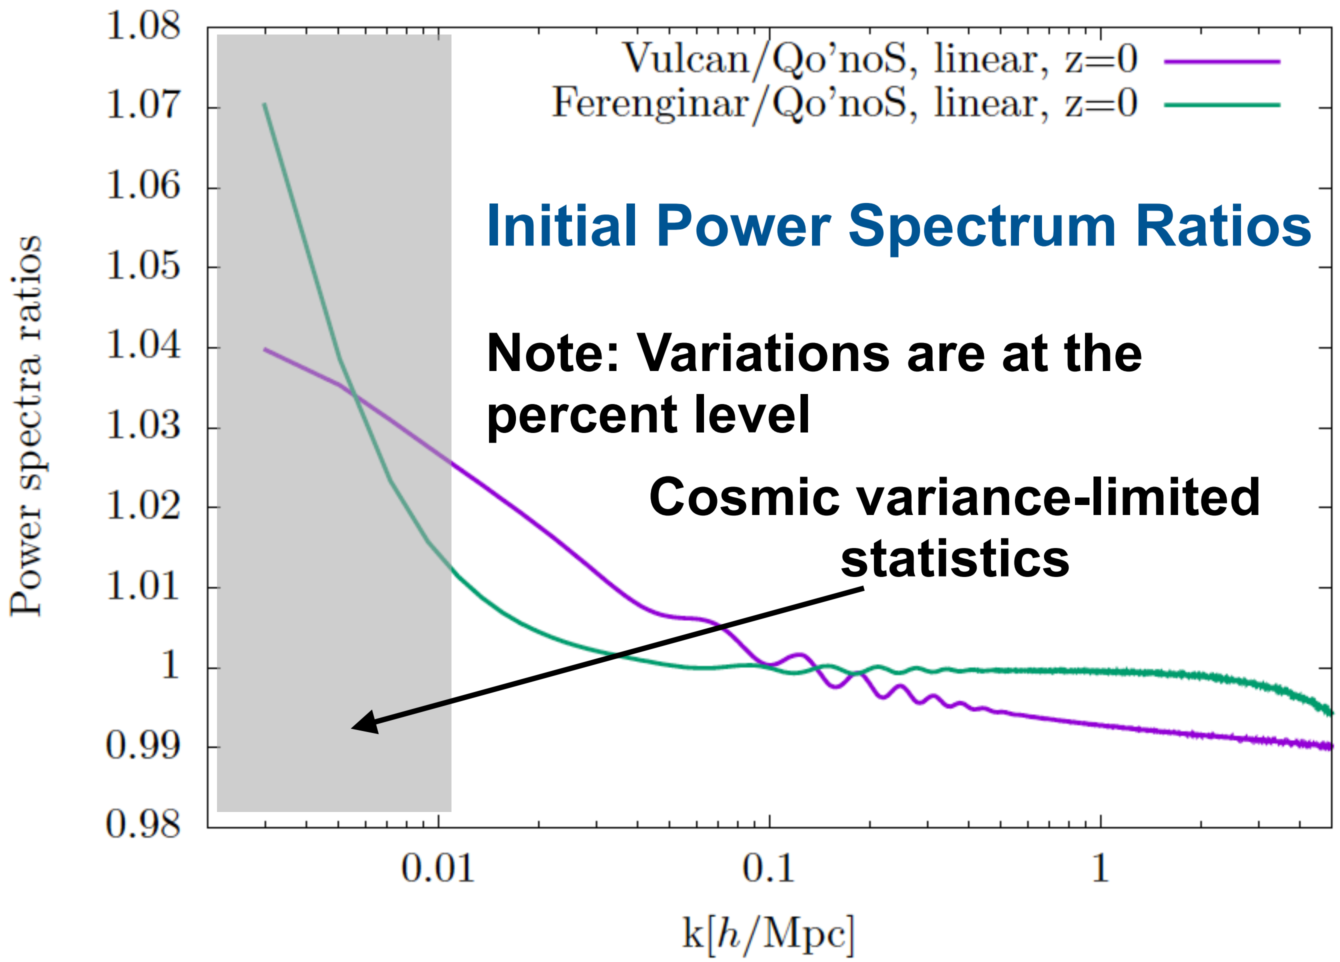

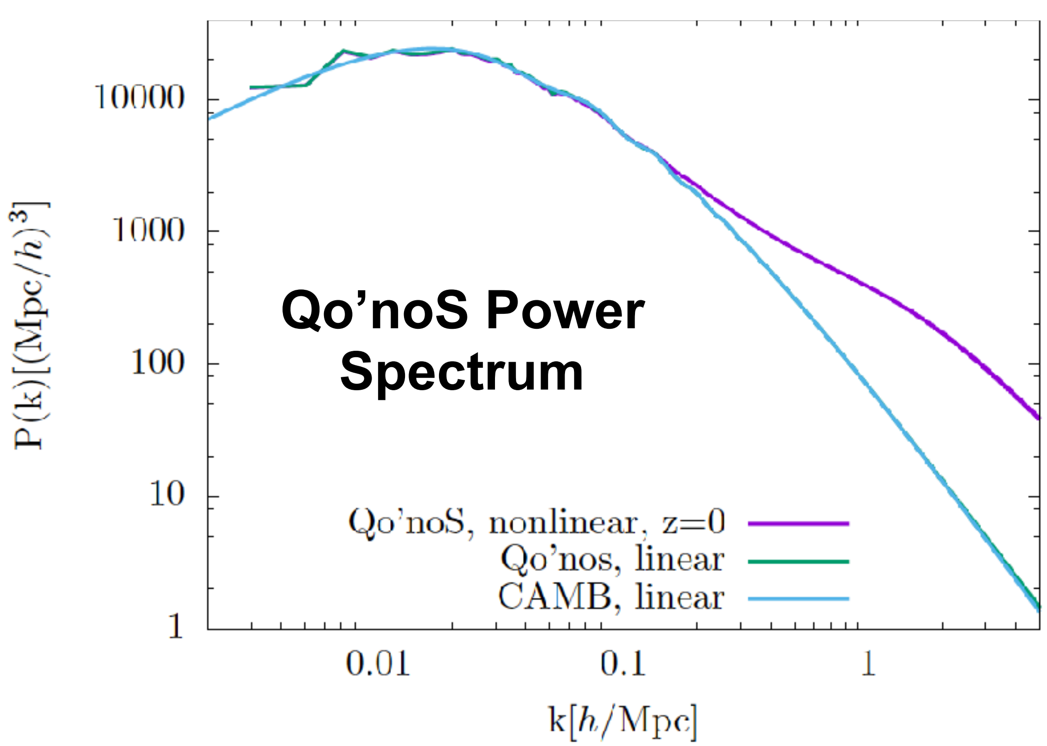

Final Power Spectrum Ratios

Nonlinear scales can distinguish subtle cosmological affects, these effects are large enough to measure with surveys, but need to worry about 'baryonic effects'

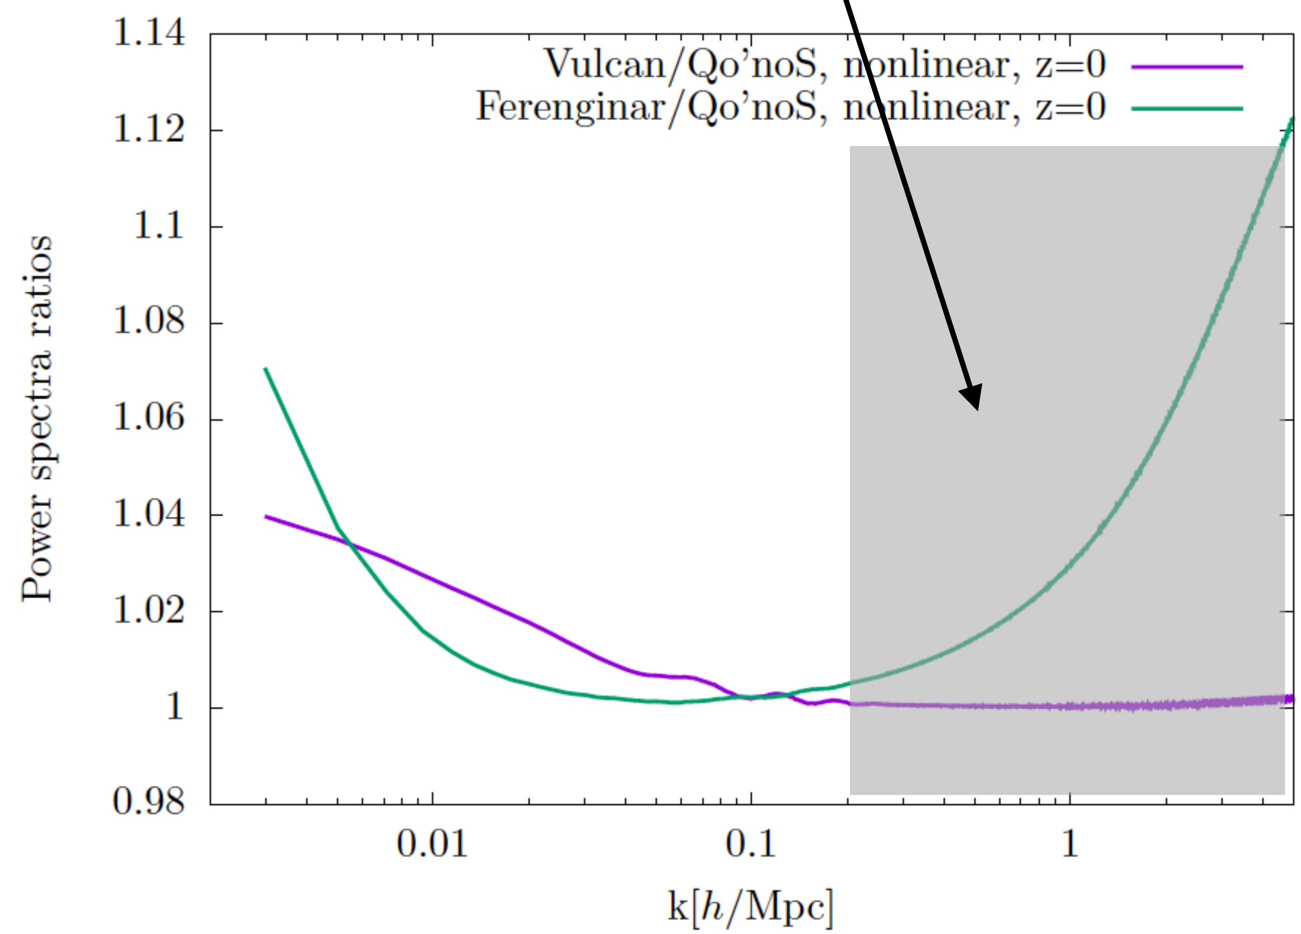

HACC on Summit IVa

- **Halo Characterization:** Jeans instability in an expanding Universe eventually forms a complex "Cosmic Web" with dark matter clumps, termed "halos", at the nodal points
- Cosmological significance of halos
  - Halos form hierarchically through mergers
  - Halo statistics are probes of the deeply nonlinear regime of structure formation
  - Galaxies form within halos
  - Galaxy properties are strongly correlated with those of their host halos and the halo's dynamical history
  - The halo mass profiles can be directly measured (and also in cross-correlation)
  - Extreme-scale simulations resolve individual halos and allow for realistic statistical samples (billions of halos)

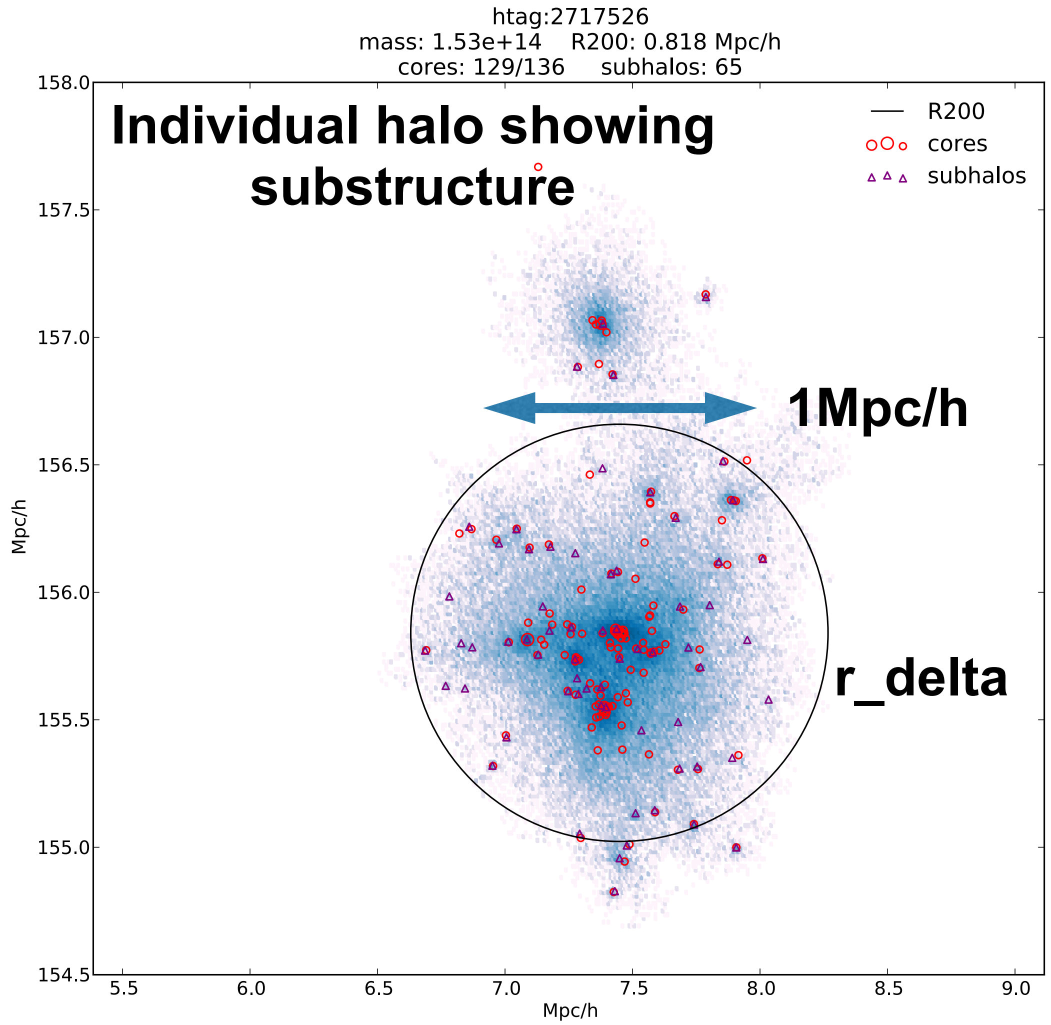

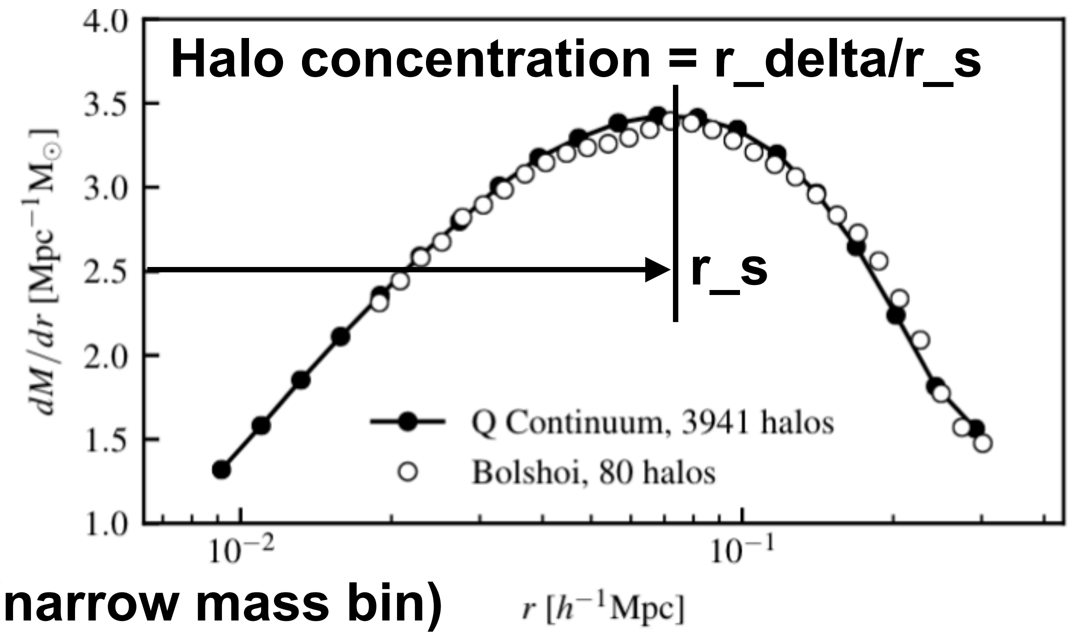

Halo profile (narrow mass bin)

HACC on Summit IVb

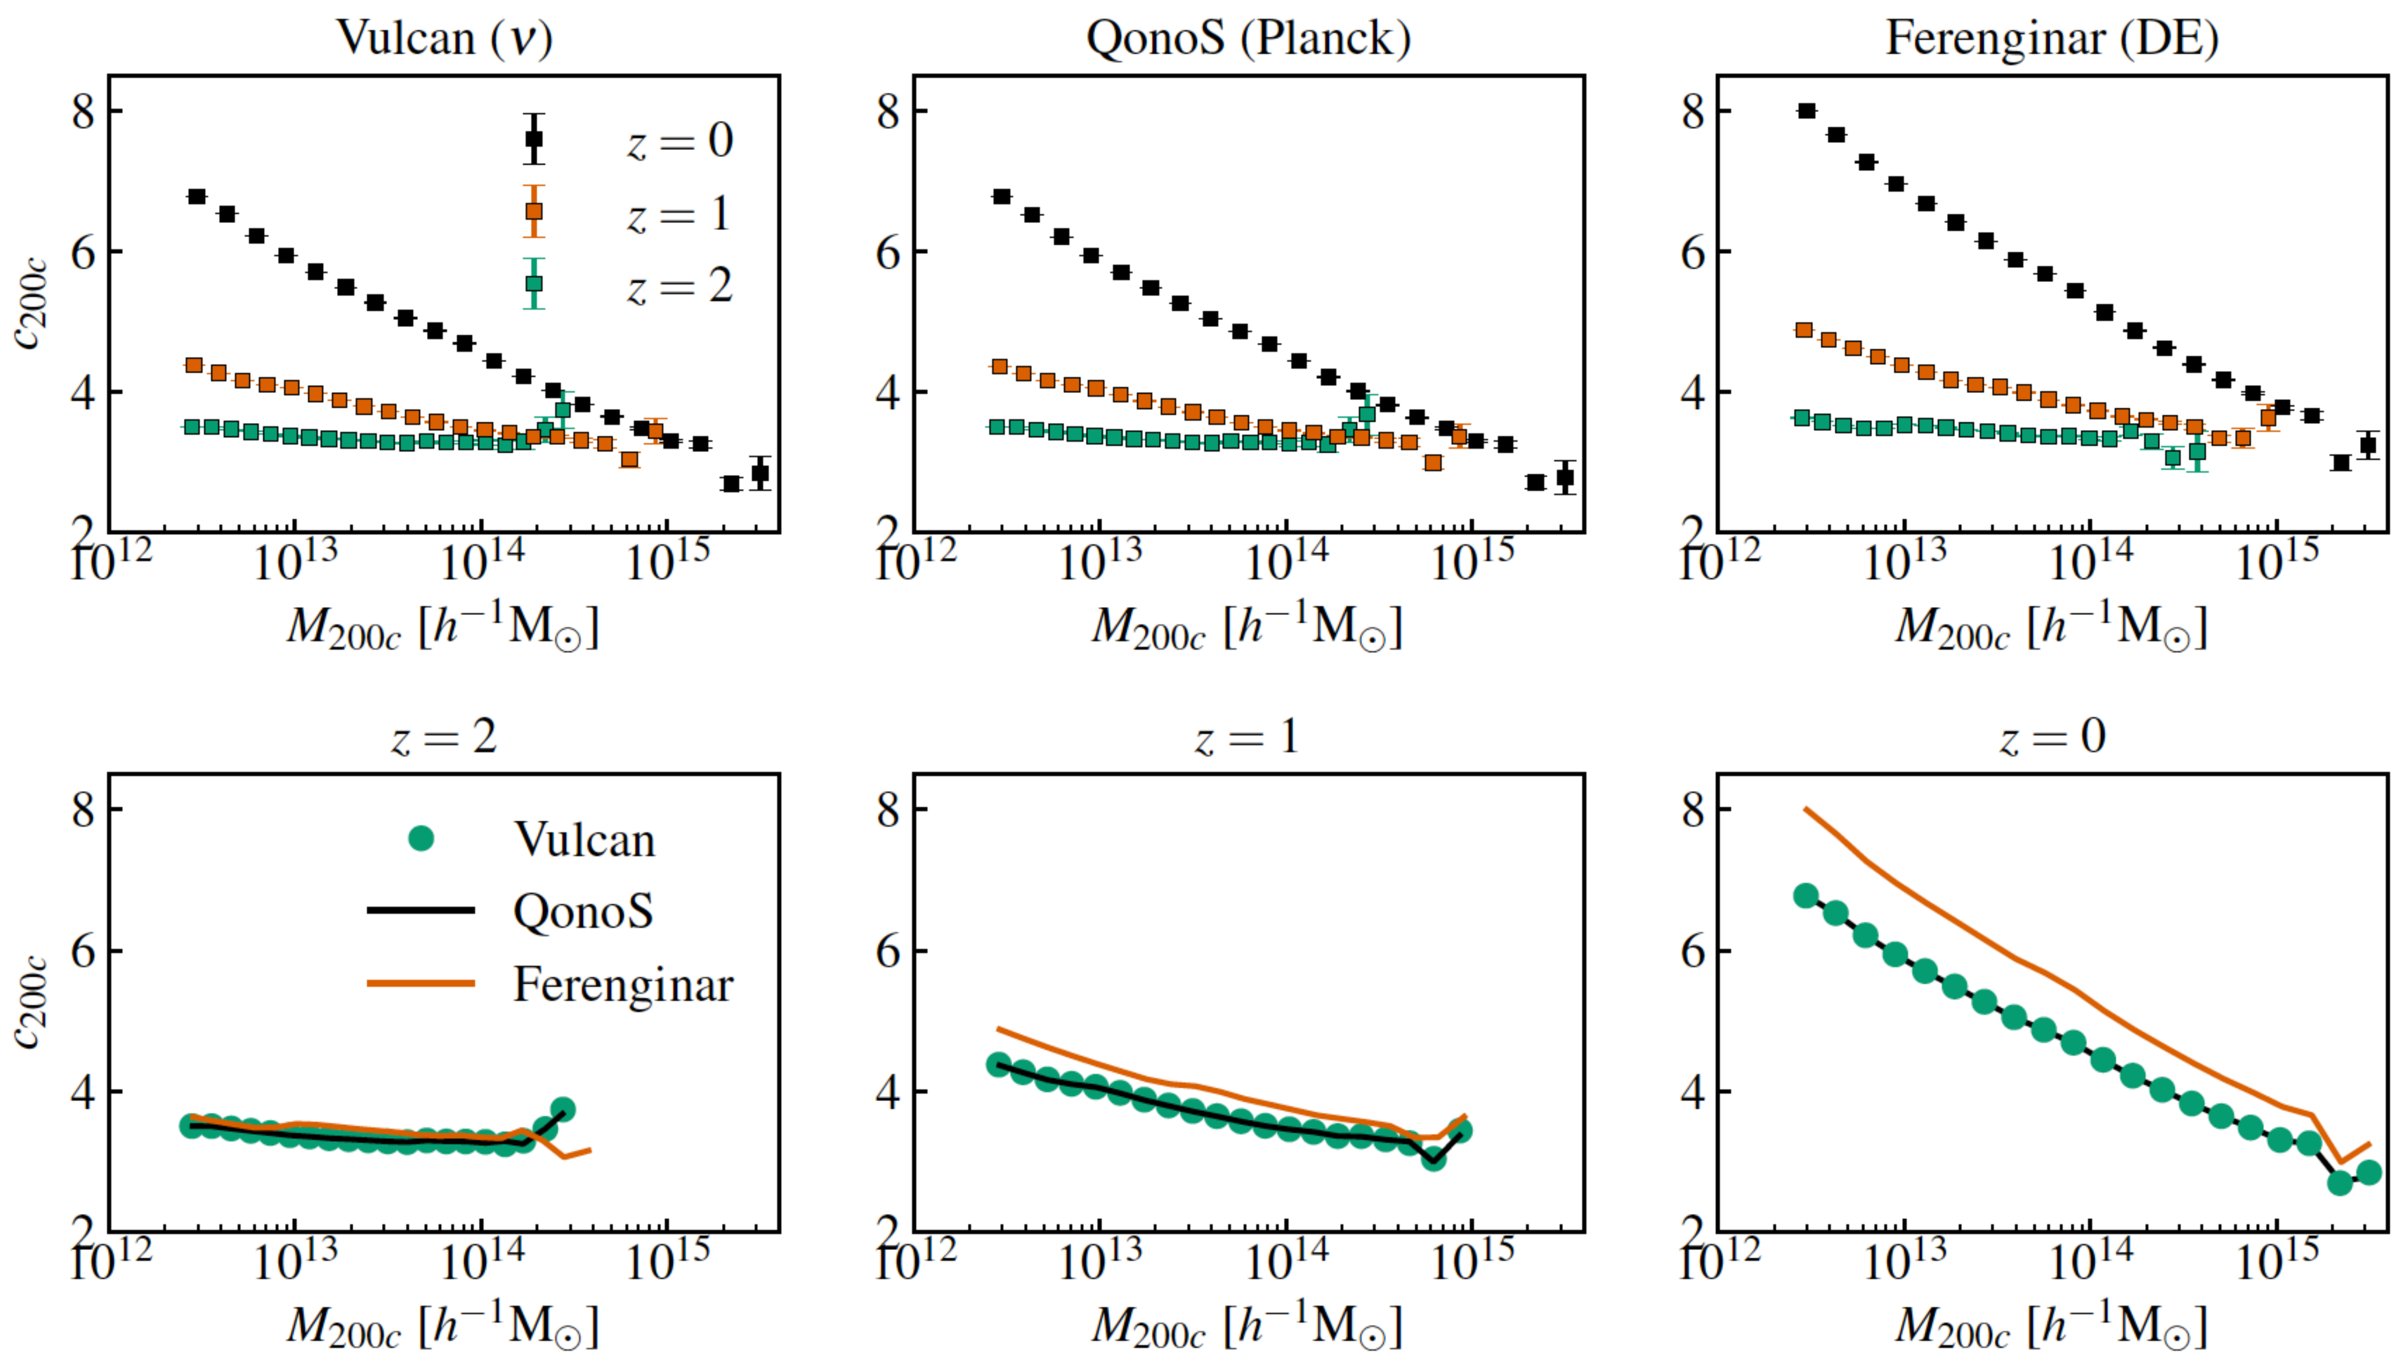

High mass resolution and very large volumes yield excellent halo statistics, here shown using the concentration-mass (c-M) relation for the different cosmologies. Note that the early dark energy model leaves a clear imprint on the c-M relation, but massive neutrinos have almost no effect since they can't cluster at small scales.

Summary

- **Ramifications:**
  - In situ analysis became the bottleneck
  - 12,288^3 FFTs were fast enough but will be sped up (~10secs/FFT)
  - Excellent IO performance with HACC’s GenericIO, few minutes per checkpoint without optimization
  - On the fly lightcone generation for the future

- **What’s Coming:**
  - Cosmological hydrodynamics
  - Optimizing CRK-HACC for Summit (successfully tested for scaling already)
  - Working on a new set of subgrid models (gas cooling, UV heating, star formation, Sn and AGN feedback, —)

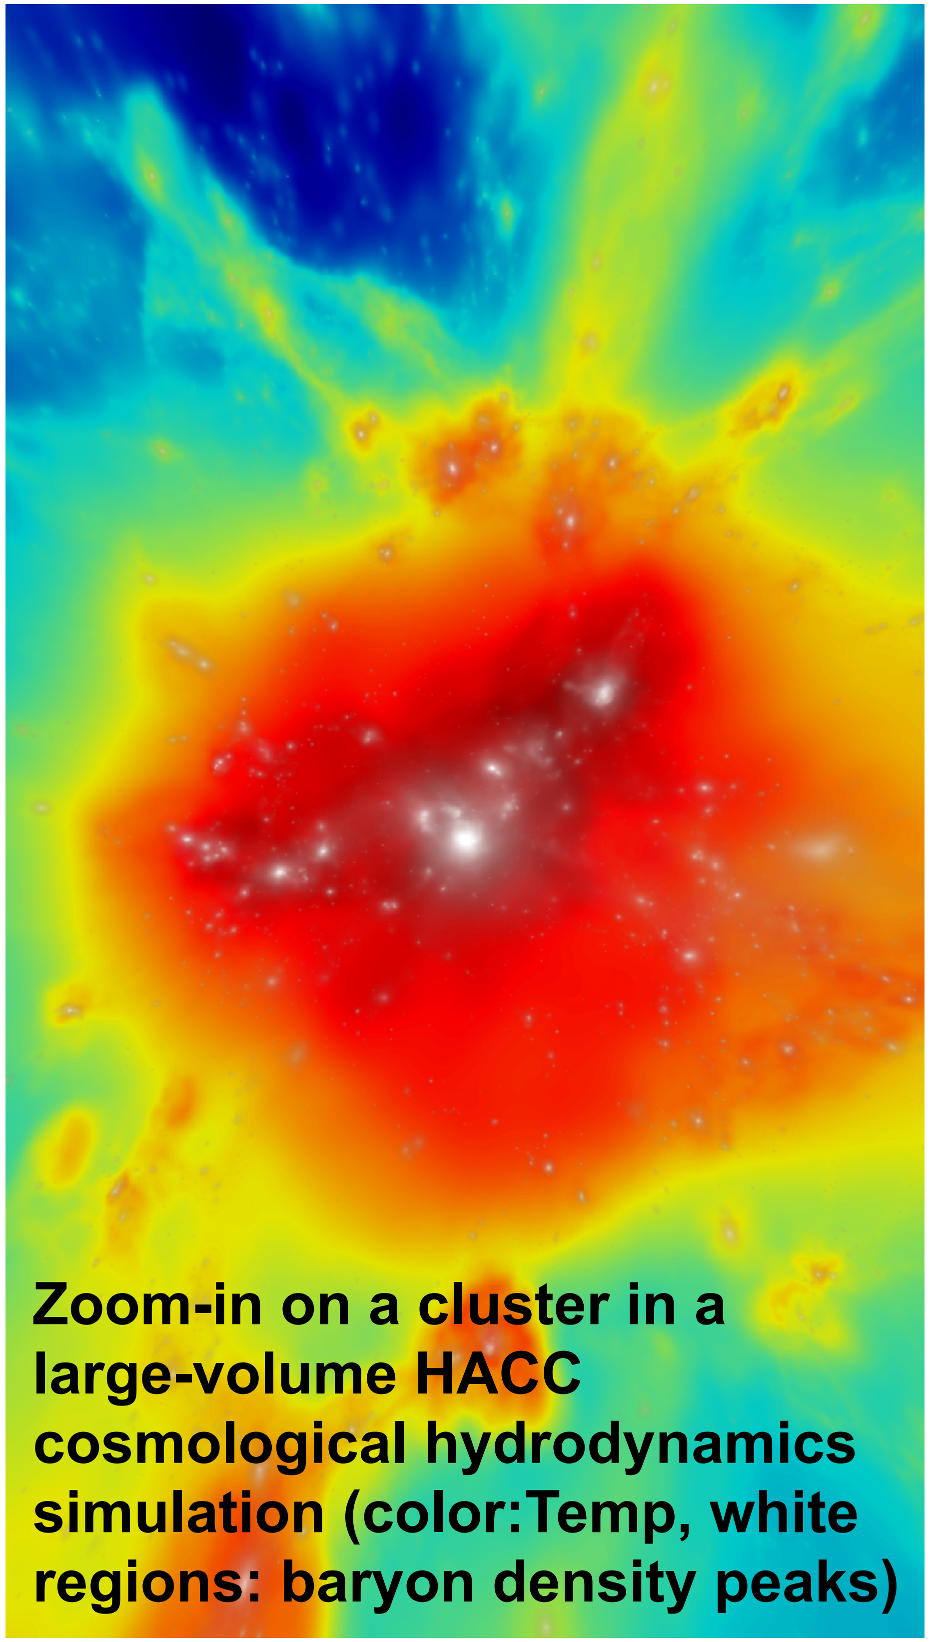

Zoom-in on a cluster in a large-volume HACC cosmological hydrodynamics simulation (color:Temp, white regions: baryon density peaks)